TRANSPORTATION ANALYSIS SIMULATION SYSTEM (TRANSIMS)

Version: **TRANSIMS-LANL-1.0**

VOLUME 0 – OVERVIEW

May 28,1999

LA-UR 99-1658

TRANSIMS

Version TRANSIMS-LANL-1.0

VOLUME 0 – OVERVIEW

May 28, 1999

LA-UR-99-1658

The following persons contributed to this document:
C. L. Barrett*
R. J. Beckman*
K. P. Berkbigler*
K. R. Bisset*
B. W. Bush*
S. Eubank*
J. M. Hurford*
G. Konjevod*
D. A. Kubicek*
M. V. Marathe*
J. D. Morgeson*
M. Rickert*
P. R. Romero*
L. L. Smith*
M. P. Speckman**
P. L. Speckman**
P. E. Stretz*
G. L. Thayer*
M. D. Williams*

* Los Alamos National Laboratory, Los Alamos, NM 87545
** National Institute of Statistical Sciences, Research Triangle Park, NC

Acknowledgments

This work was supported by the U. S. Department of Transportation (Assistant Secretary for Transportation Policy, Federal Highway Administration, Federal Transit Administration), the U. S. Environmental Protection Agency, and the U. S. Department of Energy as part of the Travel Model Improvement Program.

CONTENTS

1. INTRODUCTION

1.1 Purpose

This document provides a high-level, conceptual overview of the TRANSIMS software package. It explains the end-to-end capability of TRANSIMS by describing the major data that must be prepared, the major software methods that use these data, and the major output data produced by the software. The document is intended as a primer for potential end users and potential software vendors who might package the system for long-term use. This document does not provide a detailed description of individual data items or data structures, nor does it provide a detailed algorithmic description of the software methods. These more-detailed descriptions are provided via the Internet at *http://transims.tsasa.lanl.gov*.

TRANSIMS has flexible data requirements and modular software methods, along with multiple hardware and software operating environments. All of the modules in TRANSIMS are uniformly based on the representation of individual synthetic travelers. Accordingly, TRANSIMS is more data- and computation-intensive than existing methods. However, it is possible to begin the use of TRANSIMS by converting existing data to TRANSIMS formats and using minimally configured hardware. Some existing issues can be adequately addressed in this manner. Most issues, however, will demand more extensive data sets and more powerful computing environments to realize the full power of TRANSIMS. These needs are described in the next section. Over time, TRANSIMS modules and methods within those modules can be replaced by new, more powerful methods developed by users and researchers. These new methods will work well with other TRANSIMS modules provided that the framework data and communication protocols connecting modules and methods are followed.

1.2 Background

The goal of the TRANSIMS project has been to conduct major research and development of fundamentally new approaches to travel forecasting. TRANSIMS emphasizes and quantifies aspects of (1) activity-based travel demand, (2) intermodal trip planning, (3) traffic microsimulation, and (4) air quality and other macro analyses within the single, unified architecture shown in Figure 1.

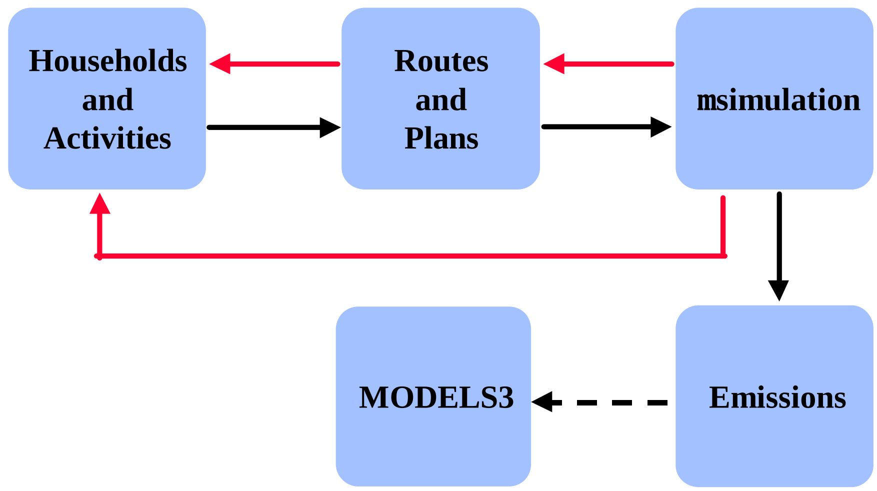

Figure 1 The highest level view of TRANSIMS consists of four major modules: Activity Generator, Route Planner, Traffic Microsimulator, and Emissions Estimator. Feedback based on the modules is used to re-plan and modify activity demand and as a modeling tool. Multiple analyses, in addition to air quality analyses, can be performed using microsimulation results.

TRANSIMS departs from traditional demand forecasting and impact analysis methods commonly in use. New technical approaches in TRANSIMS respond to issues derived from legislation such as the Intermodal Surface Transportation Efficiency Act and the Clean Air Act. Transportation planning issues requiring new technical approaches include (1) congestion pricing, (2) alternative development patterns, (3) transportation control measures, and (4) motor vehicle emissions. In addition, major programs such as the Intelligent Transportation System require new TRANSIMS analytical approaches for substantive evaluation of their effectiveness.

A major TRANSIMS technical feature is that the identity of individual synthetic travelers is maintained throughout the entire simulation and analysis architecture. All synthetic travelers are generated as part of a synthetic population developed for a specific metropolitan region using a variety of data sources including census, surveys, etc. Activity times and locations are computed for each individual. The plans generated by the Route Planner maintain individual identities, as does the Traffic Microsimulator. The resulting simulation output can provide a detailed, second-by-second history of every traveler in the system over a 24-hour day. A variety of impact analyses can be conducted using these results. This approach relies on a simple, consistent architecture—one that provides planners with deeper insight into the underlying, second-by-second dynamics of the traffic system under different local (e.g., traffic signals) and global (e.g., congestion) conditions.

TRANSIMS consists of a series of building blocks or modules that produce populations, activities for populations, routes for travelers, and microsimulated traffic dynamics. The TRANSIMS framework allows these modules to be executed in any desired order by a set of selector scripts. As each module is executed using the framework with the selector scripts, chosen data is collected in an *iteration database* to be used by the other modules. This allows for information from, say, the microsimulated traffic dynamics to be fed back to the route generator to produce a new set of more realistic routes for selected travelers.

Feedback and the control of feedback is an important function of the framework. One of the functions of feedback is to stabilize the results. As in assignment methodologies, the Route Planner may place more vehicles on links than the capacity of the link may allow. This may cause congestion to spill back onto other links. Results of the microsimulation in these cases can be fed back to reroute selected travelers to stabilize this situation. The other function of feedback is to model various aspects of transportation systems such as Intelligent Transportation Systems (ITS) implementations. Many ITS applications involve the movement of traffic information to selected travelers. Feedback accomplishes this.

The modules of TRANSIMS can be easily replaced or modified without redoing the entire TRANSIMS framework. Additionally, new modules may be added without much trouble. For example, a refined vehicle ownership module could be created to assign vehicles to an already created synthetic population without any change in the TRANSIMS framework.

2. Framework Description

The TRANSIMS framework is shown in more detail in Figure 2 and Figure 3.

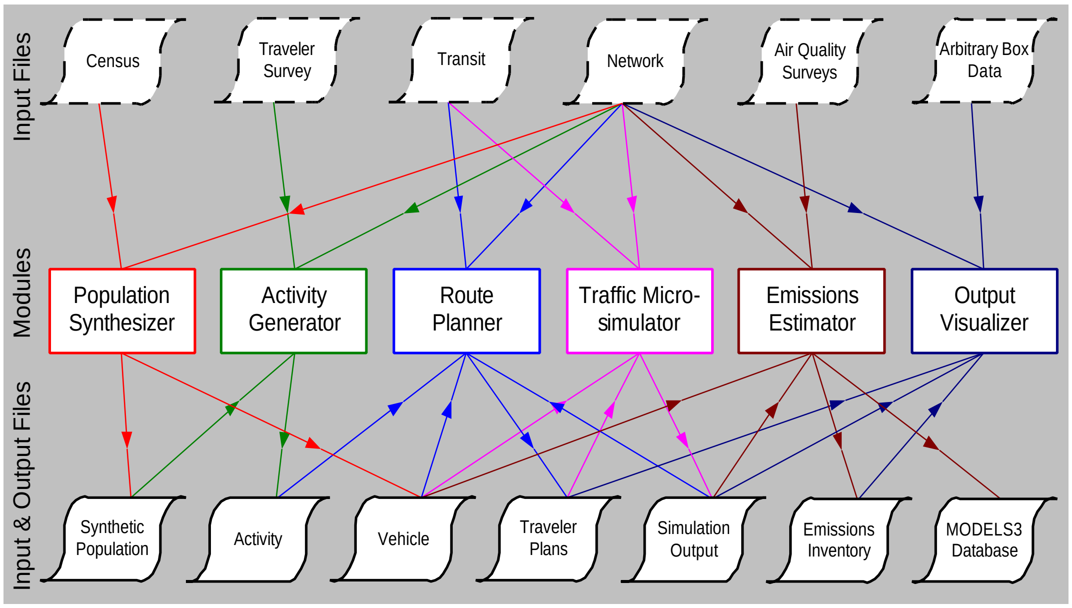

Figure 2: The TRANSIMS architecture from the perspective of *data flow*. The major TRANSIMS modules are represented in the middle row as boxes. Each of the TRANSIMS modules depends on external data that are shown in the top row. Data produced by the modules, depicted along the bottom row, are used as input to other modules. The flow of data is controlled by the framework as shown in Figure 3.

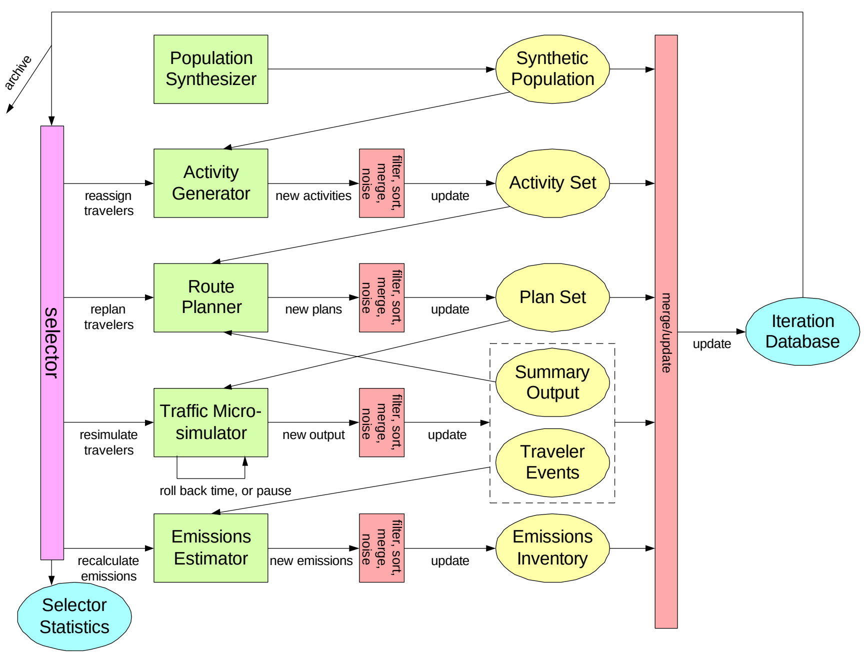

Figure 3: TRANSIMS framework diagram with the Selector at the left.

This section provides an overview of the framework by emphasizing the flow of data and computation. From left to right, Figure 2 focuses on a single pass, *feed forward* data flow to simplify the explanation of the process. This flow is controlled by the framework and a series of *selector* scripts. Figure 3 describes the framework with the Selector component. The framework is the series of modules and their file interfaces. The Selector controls the flow of information among the modules. This flow may be as simple as the *feed-forward* depicted in Figure 2, or it could be as complex as a series of feedback loops that takes data that have been produced by one of the modules and stored in the iteration database and selectively feeds them back for iterative computation by the Activity Generator and Route Planner modules.

A single feed forward pass through the TRANSIMS framework, as shown in Figure 2, begins with the preparation of a synthetic population data file for the metropolitan area. Although a variety of data sources can be used as input to the Population Synthesizer, to date we have used the following census data as sources:

1) U.S. Census Bureau Public Use Microdata Samples (PUMS)

2) U.S. Census Bureau STF-3A data

3) MABLE/GEOCORR data

The first two sources are available on CD-ROM from the Census Bureau; MABLE/GEOCORR is available over the Internet at the following web site location: *http://plue.sedac.ciesin.org/plue/geocorr*. The Population Synthesizer also uses land use data to locate households relative to the transportation network.

The second data input required is the TRANSIMS networks. In contrast to most existing network representations, network data requirements for TRANSIMS are more detailed and contain more data. TRANSIMS networks contain information similar to a TRANSYT-7F network requirements and include features such as the number of lanes, presence of turn pockets and merge lanes, lane-use restrictions, high-occupancy-vehicle lanes, turn prohibitions, and speed limits. Data on the location and type of signalized intersections are also required to produce realistic traffic flows.

The TRANSIMS networks may be viewed as "layered," where each mode is represented by a layer of the network. TRANSIMS allows walking as a travel mode (the Route Planner determines appropriate time delays for walking, but pedestrian movement is not microsimulated). Accordingly, a network especially for walking may be included in the TRANSIMS networks data file. A useful, quick way to develop a walking network is to consider it as a subset of the roadway network. Transit data are incorporated when appropriate. The information contained in this file provides the routes, operating schedules, and capacities for each transit asset (e.g., a bus or light rail train). Each link on the vehicle network may be assigned zero to many parking locations. Finally, zero to multiple activity locations may be assigned to links of any type. Land use data is associated with the activity locations. These data are used to determine the locations of specific activities.

Not all traffic in a metropolitan area is caused by residents of that area. Non-Residential Travel data include information on freight movements and itinerant travelers (i.e., those travelers passing through the metropolitan region whose trips originate outside the region). Currently, this information is derived by using the same methods that generate trip tables for these types of trips in the four-step process.

The Activity Generator module takes as major input the households in the synthetic population, local area surveys, Non-Residential travel data, TRANSIMS networks, and land use data. The Activity Generator produces a list of activities for each traveler in the system and for each freight-hauling truck. For travelers contained in the synthetic population, activity patterns and mode choice preferences are derived from surveys. This derivation depends on demographic information contained in the synthetic households. Locations for the activities of each traveler are currently chosen by using methods derived from gravity models.

The Route Planner module attempts to produce plans for every individual and freight shipment listed in the activities list. The Route Planner computes a shortest or least-cost path for each traveler. If mode preferences are recorded for the traveler, the Route Planner ensures that these are met and that the plan contains the required modal legs. The Route Planner estimates the time that it takes to make a trip based on link traversal time estimates contained in the TRANSIMS networks or in simulation output. The Route Planner uses household vehicle information derived from the synthetic population data file and transit service from the transit routes, stops, and schedules.

The Traffic Microsimulator module simulates the travel of individual vehicles and travelers in accordance with the plans provided by the Route Planner. Each plan has a specified start time, which begins the execution of movement for that traveler. Plans that overlap in time are executed simultaneously by the Traffic Microsimulator. The interactions of travelers and vehicles on the various networks over time create traffic flow dynamics. A one-second update interval ensures that dynamic vehicle behaviors (e.g., acceleration, deceleration, mode transfer, and signal intersection behavior) are captured with enough fidelity to generate realistic overall traffic behavior.

The Emissions Estimator module uses results from the microsimulation to predict tailpipe emissions for light- and heavy-duty vehicles. Pollution from spilled fuel evaporation is also estimated. These emissions are aggregated to provide input to the MODELS-3 system to produce overall Regional Air Quality estimates.

The Selector component, shown in Figure 3, is the primary mechanism used to achieve internal consistency (i.e., to achieve a reasonable agreement among the travel demands expressed in the activities list, the travel plans to meet these demands, and the execution of the plans in the microsimulation) among the various computational modules. It selectively feeds back information from one module to another. For example, results from the Traffic Microsimulator may be fed back to the Activity Generator or the Route Planner. In effect, this information is used to modify some designated subset of activities and/or plans to achieve realistic overall traffic results.

Different versions of each module have been developed during the research process. The differences range from minor changes in factors or values to completely different techniques. TRANSIMS is designed to accommodate and encourage the use of different modules both during the research process and in later commercial versions. This design, or Framework, will facilitate the development and use of new modules and, ultimately, a stronger modeling package. An example of two completely different techniques for activity generation is discussed in Section 3.2.2.

3. SOFTWARE MODULES

3.1 Population Synthesizer Module

The Population Synthesizer module builds virtual households for a given metropolitan area. These households are statistically derived from and consistent with available baseline information such as standard census data and public use microdata samples.

3.1.1 Population Synthesizer Data Input/Output

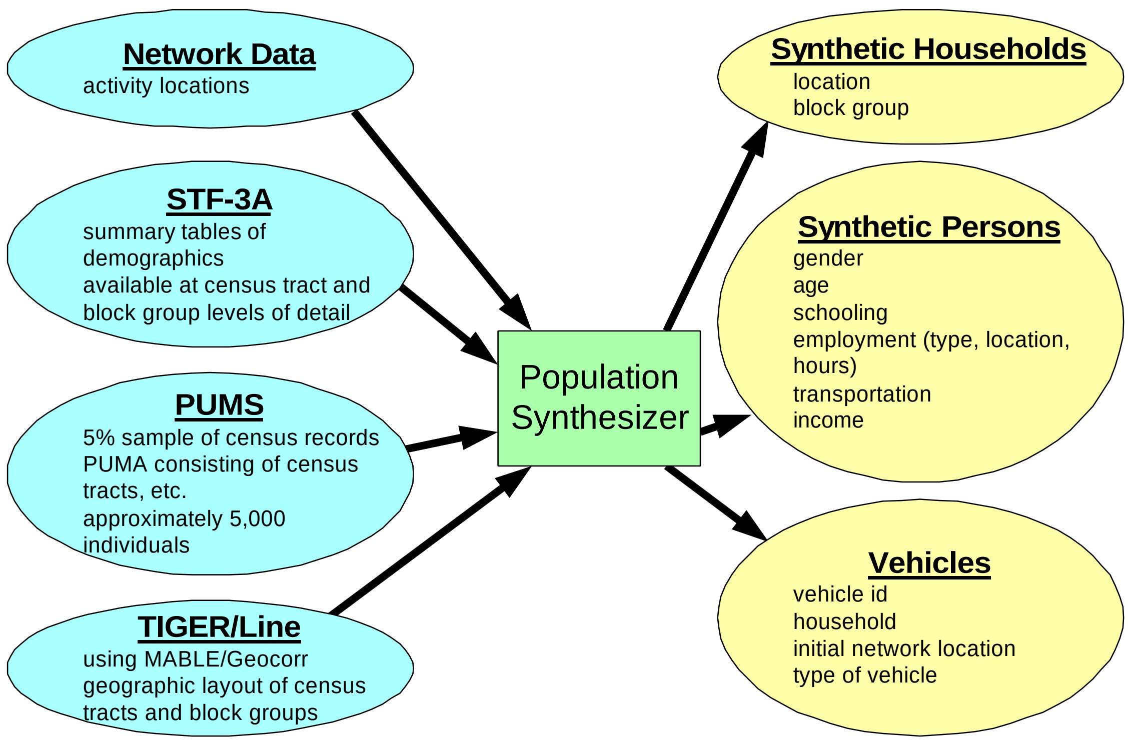

Figure 4: The Population Synthesizer uses the following source data: (1) U.S. Census Bureau Public Use Microdata Samples (PUMS); (2) U.S. Census Bureau STF-3A data; (3) MABLE/GEOCORR; (4) land use data; and (5) TRANSIMS networks. The first two sources are available on CD-ROM from the Census Bureau; MABLE/GEOCORR is available over the Internet at the following Web site location: *http://plue.sedac.ciesin.org/plue/geocorr*. The major output of the modules is a synthetic population of households containing a set of information associated with each household and household locations within the TRANSIMS networks. Forecast populations used for planning studies must contain information similar to Census STF-3A data.

3.1.2 Population Synthesizer Module Description

Any viable source of household and demographic data can be used to construct a synthetic population, provided that the output of the computation is formatted in accordance with TRANSIMS data formats for synthetic populations (see TRANSIMS, Volume 3—*Files* at *http://transims.lanl.gov*. A separately released software package called *TRANSIMS: Synthetic Population* is available. This package allows users to experiment with and build synthetic populations using census data.

Synthetic populations contain households comprised of one or more individuals. Each individual has an associated set of demographic variables (e.g., age, sex, income, etc.) Households are usually generated for small areas such as census block groups or census tracts. Households are assigned to activity locations on a link of the TRANSIMS network according to land use characteristics associated with the activity locations on that link. If a *walk* network is used, it is useful to assign households to activity locations on *walk* links only. It should be noted that multiple activity locations can be assigned to one link. In particular, an activity location on a link can represent one side of a street, while a second could represent the activities that take place on the other side; activity locations might even represent individual buildings on a street. Synthetic populations provide input to the Activity Generator. Information contained in synthetic populations can also be used to categorize and filter subsets of the population used for various types of equity analyses.

In the methodology presented here, each household in a synthetic population is classified as either family, non-family, or individuals living in group quarters such as dorms. Family households contain one or more adults and possibly children. Demographics associated with members of households vary in accordance with source data and study needs. At a minimum, a TRANSIMS household must contain at least one person.

An example of how to generate a synthetic population is outlined in the paper by Beckman, Baggerly, and McKay, 1996 [1]. This technique uses the 1990 census data given in Census Standard Tape File 3A (STF-3A) and PUMS. STF-3A gives demographic summary data on a census tract or census block group level. The PUMS data is comprised of complete census records, without any identification data such as name, address, census block, block group, or tract. PUMS covers 5% of the households surveyed over a large number of block groups. The area over which the PUMS ranges is called the Public Use Microdata Area (PUMA). The block groups that compose a PUMA are obtained from MABLE/GEOCORR, a geographic correspondence engine available at web site *http://plue.sedac.ciesin.org/plue/geocorr*.

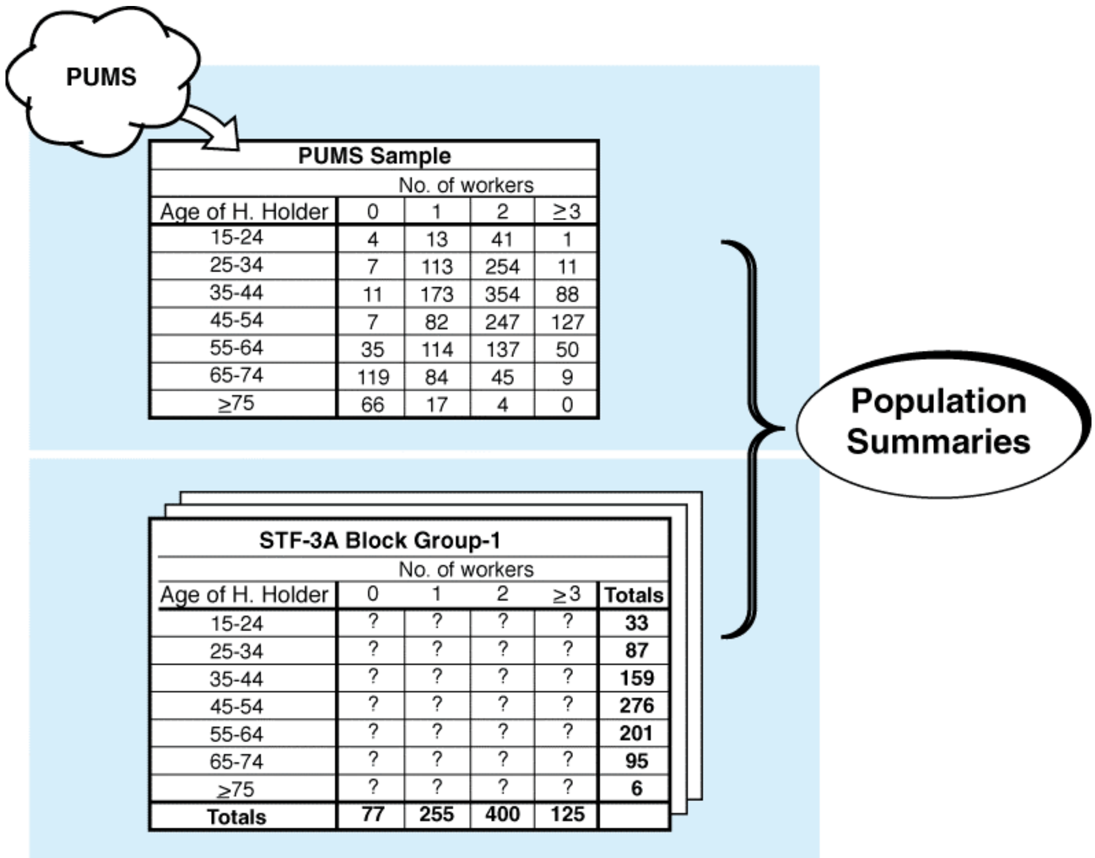

PUMS Sample

Age of H. Holder	No. of workers 0	1	2	≥3
15-24	4	13	41	1
25-34	7	113	254	11
35-44	11	173	354	88
45-54	7	82	247	127
55-64	35	114	137	50
65-74	119	84	45	9
≥75	66	17	4	0

STF-3A Block Group-1

Age of H. Holder	No. of workers 0	1	2	≥3	Totals
15-24	?	?	?	?	**33**
25-34	?	?	?	?	**87**
35-44	?	?	?	?	**159**
45-54	?	?	?	?	**276**
55-64	?	?	?	?	**201**
65-74	?	?	?	?	**95**
≥75	?	?	?	?	**6**
Totals	**77**	**255**	**400**	**125**	

Figure 5: The first step of constructing a synthetic population is to use PUMS data and standard census data STF-3A. PUMS contains a 5% sample of real census records that have been modified to maintain anonymity. These are used in conjunction with STF-3A summary data, which contains demographic information.

The construction of a synthetic population is accomplished in several steps. First, an appropriate PUMA is identified, then MABLE/GEOCORR is used to obtain a listing of block groups within the PUMA (see Figure 5). Next, summary statistics from STF-3A for each of the block groups identified in the PUMA are obtained. As an example, summary data might include the age of the householder, the family income, the number of workers in the family, the race of the householder, and the type of family. A multidimensional table is constructed from the PUMS data where the dimensions correspond to the summary statistics from STF-3A. In the example above, the multidimensional table would have five dimensions corresponding to five classifications: the age of the householder, the family income, the number of workers in the family, the race of the householder, and the type of family. The numbers of households in the PUMS for each classification compose the multidimensional table. The proportion of households for each classification for the block groups is unknown—they are determined by a two-stage iterative proportional fitting procedure outlined in Beckman, Baggerly, and McKay, 1996. This procedure satisfies the distributions of the STF-3A data for each block group while maintaining the correlation structure of the table constructed from the PUMS. A forecast population can be generated using the same techniques, but replacing the STF-3A data with forecast STF-3A data.

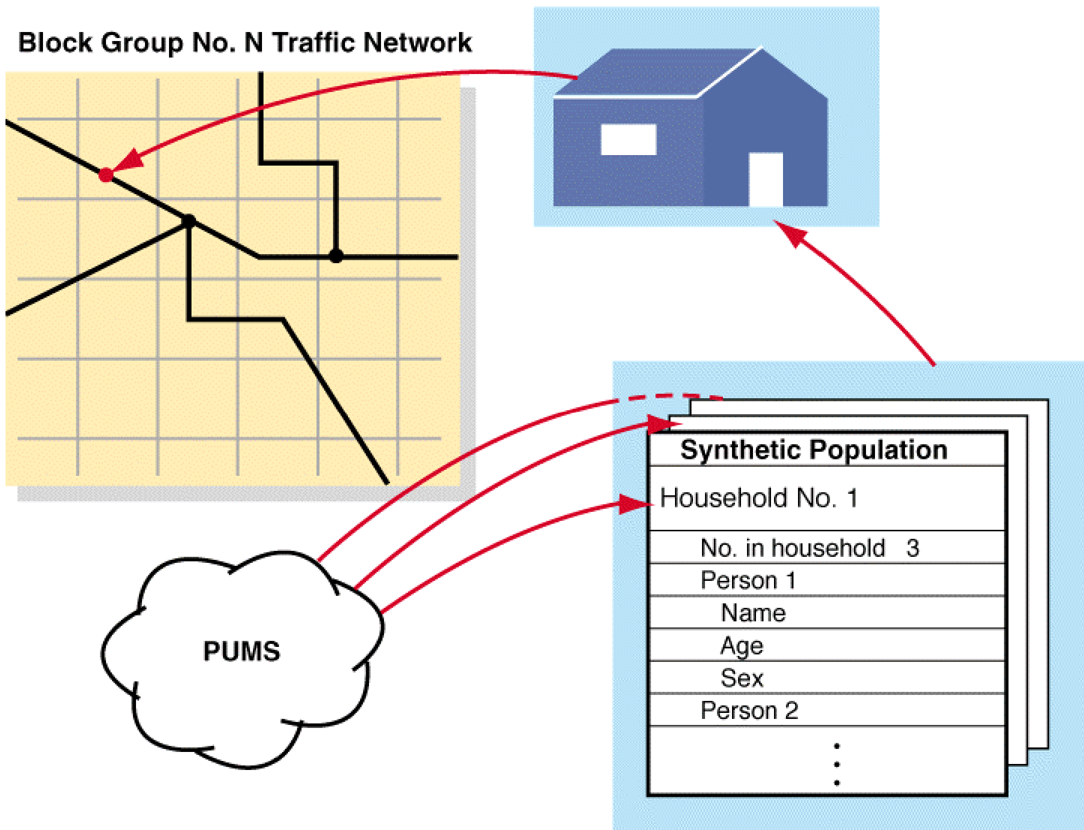

Figure 6: Creating households and placing them on the network involves randomly selecting actual PUMS households in accordance with the proportions derived from the Iterated Proportional Fitting method.

The final two steps in population generation are (1) selection of households from the PUMS to match the constructed tables, and (2) placement of the households on links in the network. An estimated proportion of households for each block group in each category of the multidimensional table is obtained from the iterative proportional fitting procedure. The synthetic population is formed for each block group by randomly selecting households from the PUMS according to these proportions.

Households are selected from the PUMS to match the number of households in the Census over a given geographic area such as a block group or a census tract. Land use information is used to place each household within a block group at an activity location on a walk link. Land use data is stored in the network activity location file. At a minimum, this file contains the identity of the activity locations, their locations, the corresponding block group and census tract, and some indication of the activities that may be performed at that location. For the Portland Metro case study, this information is represented by the area (in square feet) of land use types surrounding the link containing the activity location. Here, land use information includes the square footage of commercial, industrial, multi-family, and single-family residential property surrounding the link. Additionally, the number of employers and the total number of employees around the link are included in this file. The data in the file may be generated using aggregated data by considering the proportion of area around the link when compared to the total area represented by the aggregated data.

To place households from the synthetic population at an activity location on a link, the links within a block group are first identified. For Portland, weights for each activity location are determined by adding the residential square footage and a multiple of the multifamily square footage. Each individual weight for an activity location is divided by the total of the weights of all of the activity locations in the block group. These ratios are used as the probability of a household being located on a link. For each household, a random link (based on the probabilities) is selected, and the household is placed on that link. Note that households need not be placed at unique activity locations. Many households can share the same activity location. This household location algorithm may remain the same no matter which area is being studied. However, the weights given to the activity locations in a block group will depend on the quality and availability of land use data.

Households may be located on a link by other techniques. For example, census block data could be used to determine the number of households in a block, and this number could then be associated with an activity location on the link. Another possibility is to use electronic phone books or aerial photography to determine this number. Neither of these techniques has been used to date in TRANSIMS case studies.

Each generated household is comprised of individuals. The household and individual demographics are as rich as the data in the PUMS. Since the household is a replica of a household in the PUMS, every variable in the PUMS is also associated with the household. This includes, but is not limited to, the household structure, the individual and household incomes, individuals' races and ages, those individuals who work, and those who attend school.

At this time, the vehicle ownership is given by the number generated from the synthetic population procedure using the PUMS. A more refined vehicle ownership model based on population demographics and network characteristics could be implemented and called in the framework after the synthetic population is generated. In both cases, each vehicle is an entry in the vehicle file described in TRANSIMS Volume 3—*Files*. This file contains the TRANSIMS vehicle identification number, the household to which it is assigned, and the vehicle emissions type. The emissions type is used in the Emissions Estimator module to determine emissions. It reflects the operating condition of the vehicle, its type, and age. Presently, these vehicle types are assigned at random according to a national distribution of 23 vehicle emission types that are described in the Emissions Estimator section of this document. In the future, the analyst may wish to utilize a vehicle type model based on Department of Motor Vehicle statistics, inspection and maintenance records, and the synthetic population.

3.2 Activity Generator Module

The Activity Generator module computes a list of activities for each traveler and each freight movement. Each activity in the list has an associated activity type, location, and time. Any methodology may be used to generate the TRANSIMS activities as long as the activities are *TRANSIMS compatible*—that is, if each member of the synthetic population is assigned a set of activities and these activities are stored in an activity file in the format outlined in TRANSIMS Volume 3—*Files*. In this document, we describe two methods for generating TRANSIMS activity files. One is being developed by the National Institute of Statistical Sciences (NISS) for the Los Alamos National Laboratory, while the other is work performed for Portland Metro.

3.2.1 Activity Generator Major Input / Output

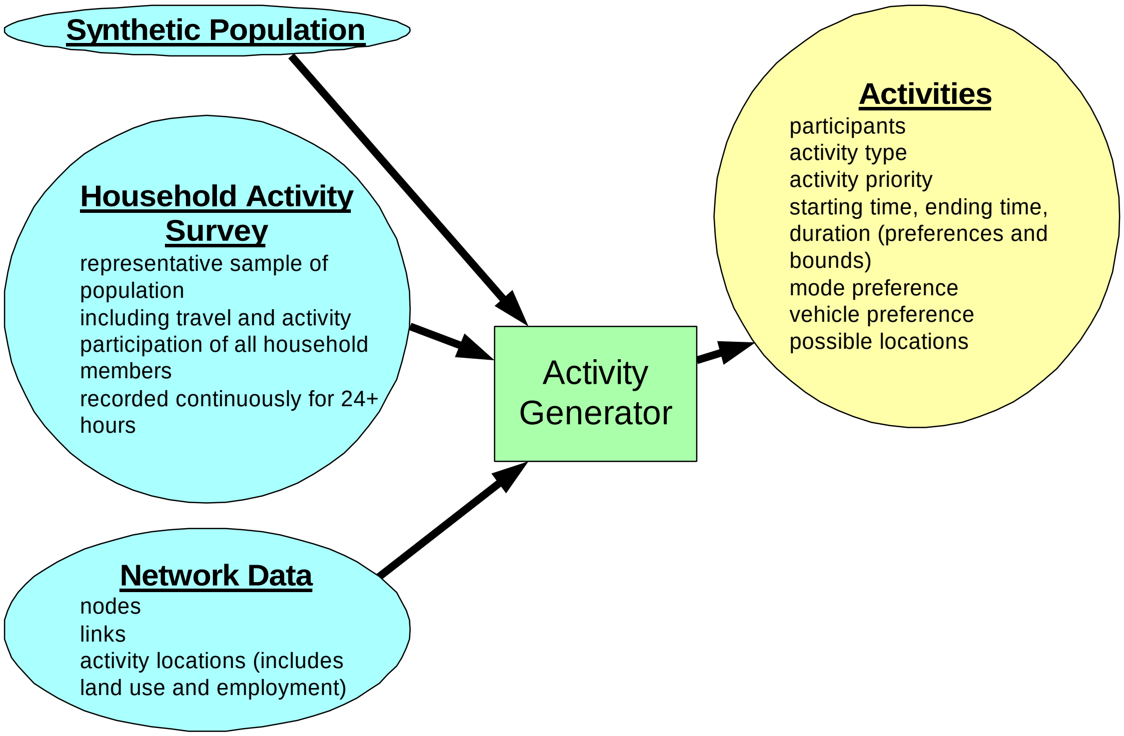

Figure 7: The Activity Generator uses the synthetic population and non-residential travel data to compute an activity list for every traveler and freight movement. This list constitutes the travel demand for the region under study. Surveys are used to determine activity patterns for households in the synthetic population.

Each TRANSIMS household member is assigned a set of activities for a 24-hour period. These activities are based on models developed from activity surveys that are collected by metropolitan planning organizations. The activity surveys generally ask respondents to list the household demographics and to report the location, type, and duration of their activities over one or a multi-day period. The survey is used to develop activity patterns for households of various types.

The output of the Activity Generator is a data file containing the activities for each individual traveler. This file identifies an individual and his/her corresponding household. Each individual is given a list of activities with the following attributes: type (e.g., home, work, shopping, school, or any other activity type collected in the survey), starting time range, ending time range, and activity duration range.

3.2.2 Activity Generator Description

TRANSIMS assumes that any two activities, separated by time and location, require travel between them. The degree of detail in both the synthetic population and activities can vary depending on the availability of data and the study being performed. A more detailed study with more realistic traffic can be accomplished with a more detailed and realistic representation of the metropolitan population and the activities in which the population engages. The only requirement is that the file structure protocols are not violated. This section outlines the structure of the activities, the methods for generating them, the minimum data required to execute TRANSIMS, and other sets of data that may be required by some TRANSIMS analyses.

Consider the following example: A person may need to arrive at work between 7:30 and 7:45 a.m. and stay until 11:45 a.m. to noon. The starting time range in the activity list would be 7:30 to 7:45; the ending time, 11:45 to 12:00 with a duration of 4 to 4.5 hours. The time information is a redundant set; consequently, not all of the fields (start time, end time, and duration) need to be completely specified. Each activity has a location specified in the file, which must be part of the TRANSIMS network data file. The location is one of the activity locations listed in the activity location network file. In addition to the time and location, a travel mode to reach the activity is also assigned. The mode may be very specific (such as requiring a car), or it may be general, allowing the Route Planner to search for the best mode(s) out of all possible modes. If two or more individuals are making the same trip, in a particular driving trip, the individuals are identified as part of the activity list.

One technique for activity generation, the NISS methodology, involves overlaying each synthetic household with a complete activity pattern. This is accomplished in the following way: An activity survey is processed to obtain the total time spent in activities by activity type for each surveyed household. These times are weighted and summed to form a measure of total time spent in activities for each household. Demographic variables of the household and the individuals in the household are chosen based on which ones make the best predictors of the activity duration time. The predictor takes the form of a decision tree where questions are asked at every level of the tree (see Figure 8).

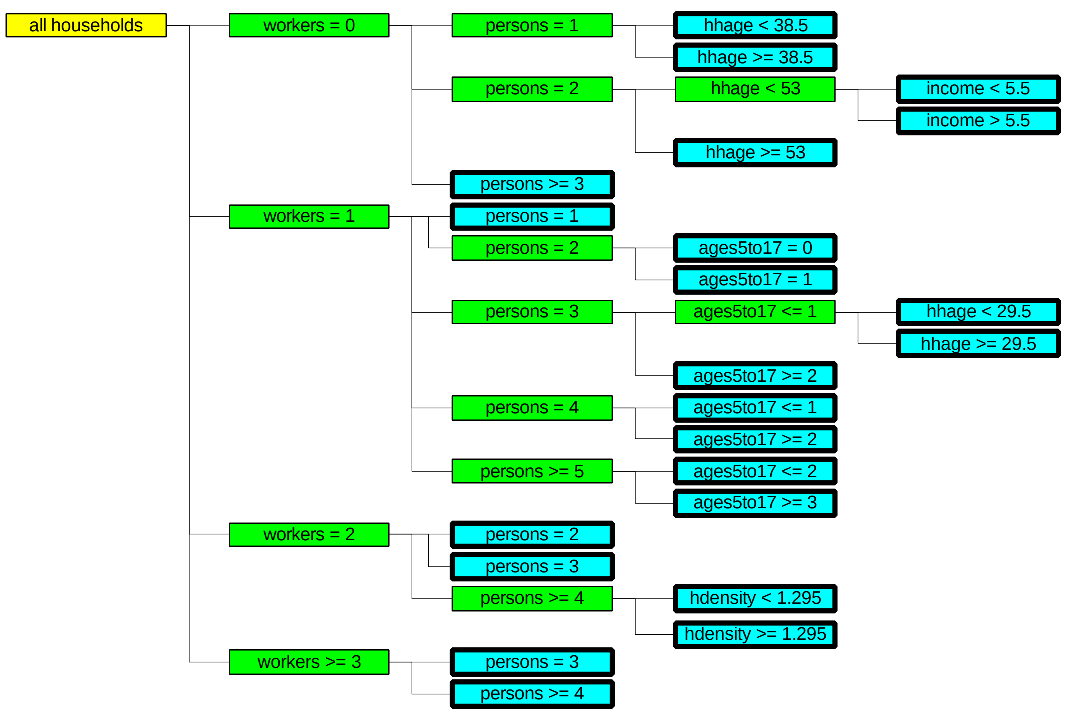

Figure 8: An example of a *decision tree* used to classify household activity patterns.

For example, the first question may be “is the household size equal to one?” The next level of question for single person households may be “is the person a worker?” while the next level of question for a multiple-person household might be “are there more than two people in this household?” The tree ends at what are called *terminal nodes,* which specify the activity pattern for the household.

Once a decision tree is constructed, each household from the survey is classified as belonging to one of the terminal nodes of the tree. More than one household is usually assigned to each of the terminal nodes. Base activity patterns are allocated to individual households in the synthetic population by classifying them according to the decision tree and giving them an activity pattern of one of the survey households that were similarly classified as belonging to that node. The activity pattern is screened and corrected for obvious errors, such as a two-year-old person working or having to drive. Given the base activity pattern, the locations of the activities are determined by a modified gravity model based on land use data in the activity location file. For Portland, these are the square footage of commercial and industrial property around activity locations on the link and the number of employees and employers assigned to the activity locations on the link. Initial activity start and end times are determined from network skim times for travel from one area to another.

A second approach to TRANSIMS compatible activity generation is being developed by Portland Metro. Using nested logit models with the parameters estimated from the Portland survey data, primary tours are assigned to individuals in the population. Some primary tours are home-to-work, work-to-home, and home-to-shop. After the primary tours are assigned to the individuals, secondary tours and intermediate stops on a primary tour are assigned to the individual by sampling the tours from the survey. Times, locations, and mode choices are obtained using gravity models.

The TRANSIMS Framework allows the activity sequence, times, and locations to be easily changed by feedback from the Route Planner and the Traffic Microsimulator. There is a tradeoff between spending resources to make the activity list as correct as possible and correcting the list by feedback. This is an active area of TRANSIMS research. Hopefully, only a rough cut is needed for the base activity list, and a refined list will easily emerge through feedback.

The importance of feedback to refine the activity list should not be underestimated. If it can be shown that feedback can improve an activity list so that extremely detailed activity models are not necessary, then national survey data and data from other sources could be used to construct the base set of activities. In this case, an extensive local survey would not be needed. Future TRANSIMS research in this area is important and may reduce the cost of surveys for metropolitan areas.

At this time, freight and itinerant travelers are handled in TRANSIMS through trip tables, which are derived from a portion of the four-step process. The lack of data and models in these two areas makes this simplifying approximation necessary. This can be changed easily as better freight models are developed. Trip tables are transformed into TRANSIMS activities in the following way. Individuals (not created as part of synthetic populations) are created as vehicle drivers, and one vehicle is created for each trip in the trip tables. There may be multiple trip tables representing various types of vehicles. For example, numerous trip tables could be used to represent trucks with different numbers of axles. A trip table must be available to accommodate the microsimulation of itinerant travelers. In this version of TRANSIMS, a single vehicle is used for only one trip. If a new model generates a sequence of trips for, say, a single delivery truck, this would be implemented easily by creating an activity list with an activity at every stop. Each truck (and bus) simulated in TRANSIMS is listed in the vehicle file. As with automobiles, trucks and buses are assigned an emissions type.

3.3 Route Planner Module

The Route Planner module develops travel plans based on the *demand* represented in the activities data file. Each traveler, including itinerant travelers, truck drivers, and transit drivers has an individual travel plan. After the plans are developed for the travelers, they are simultaneously executed in the Traffic Microsimulator.

3.3.1 Route Planner Major Input/Output

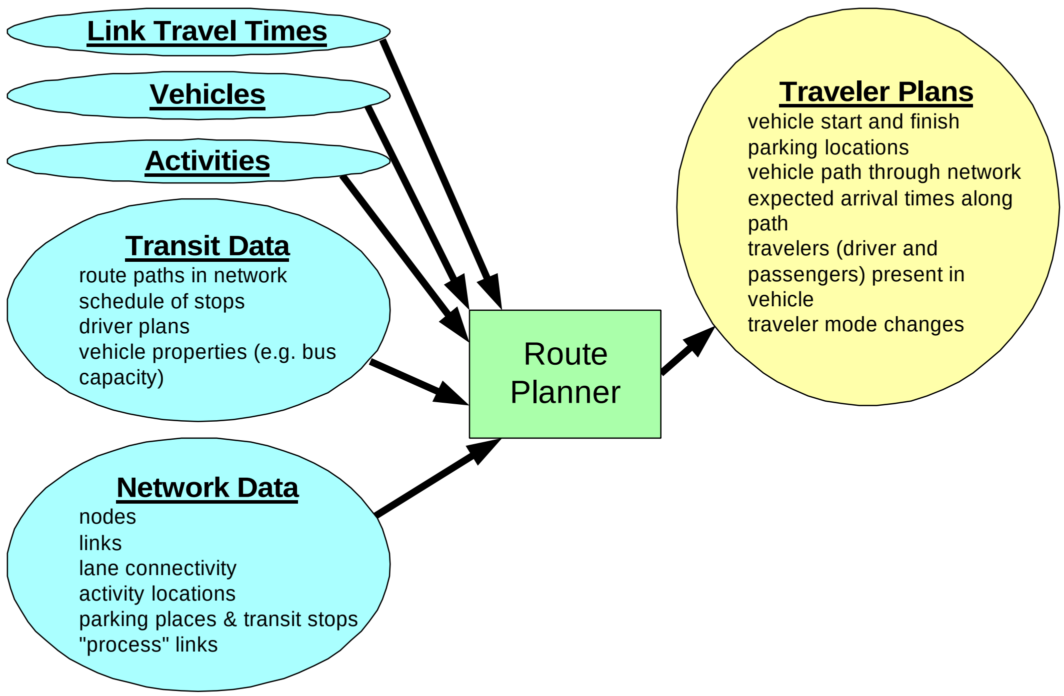

Figure 9: The major input to the Route Planner includes: (1) TRANSIMS network data, (2) activities data, (3) transit data, and (4) vehicle information from the synthetic population data.

The major inputs to the Route Planner are the activities list, the TRANSIMS multimodal network, and the vehicle file as shown in Figure 9. Travel demand is implied by the activities list when two activities for an individual take place at different times and locations. The TRANSIMS multimodal network is transformed to an intermodal planner network for routing purposes, and the initial starting position of all vehicles is maintained in the vehicle file. Since TRANSIMS tracks the movements of each individual throughout the simulation, the Route Planner retains the location of each individual's vehicle. This allows an individual to drive to a parking location, walk from the parking lot to work, then return to the same parking location to retrieve the vehicle for the trip home.

3.3.2 Route Planner Description

The Route Planner computes the *shortest path* for each traveler in the system, subject to constraints. Each link within the network system has a cost associated with it. Costs for a link can be computed simply with input such as an estimated time delay. It can also be computed in more sophisticated ways—for example, as a function of several variables including time delays and actual monetary costs of the link, or including more abstract variables such as a penalty for traveling through construction areas. Accordingly, *shortest path* can also be interpreted as *least cost*. Constraints are provided by criteria such as mode preferences for different legs of the trip.

The Route Planner has several distinguishing features:

- First, plans are computed for each individual traveler in the population, based on that individual's activity demands and preferences. This allows each traveler to have an individualized view of the transportation system. Accordingly, costs associated with links in the network are computed for each traveler.

- Second, link costs are computed in a time-dependent manner, which can account for time delays resulting from actual travel conditions such as peak-hour congestion.

- Third, the Route Planner abides by any mode preferences contained in the activity files. Thus, if the activity files specify that a traveler will first walk, then take a car, and then walk again between two desired activities, the Route Planner will produce a plan (if feasible) that ensures these modes are used in this sequence.

The Route Planner assumes the following definitions: A *traveler plan* is the set of trips that carries the traveler through his/her desired activities. Trips may be separated by activities that have a given duration. A *trip* is a set of contiguous legs. A *leg* is a set of contiguous nodes and links that are traversed with a single travel mode. For example, a trip may consist of three legs, the first walking, the second transit, and the third walking; and a traveler plan could consist of a trip from home to work, work to shopping, then a trip from shopping to home.

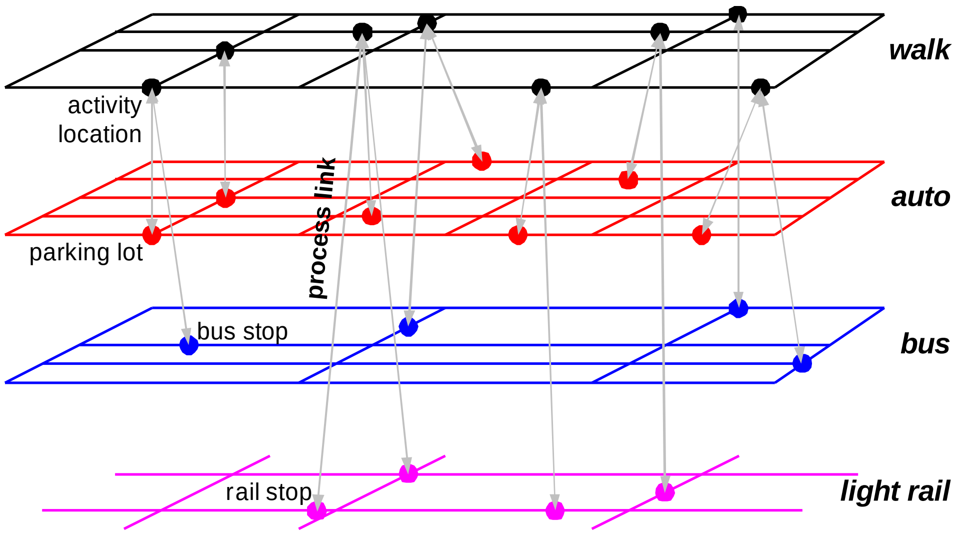

Figure 10: A high-level depiction of the various network levels used by the Route Planner. From individual traveler preferences and constraints contained in the synthetic population and activities data blocks, the Route Planner must plan for trips comprised of multiple modal legs (e.g., walk-car-walk). Constructing multiple network levels where each level can be encoded as a different modal network allows for the efficient calculation of trips constrained by modal sequences. Also shown are the process links connecting the modal networks.

As shown in Figure 10, the Route Planner conceptually views the network as a set of interconnected, unimodal layers. That is, for each mode (e.g., walking, car, transit, etc.) a separate network exists. At certain designated nodes in each of these networks, a special link, called a *process link*, exists to one or more of the other unimodal networks. These links allow intermodal transitions to occur. The process links are considered to be part of the walking network. The layered networks are constructed from the underlying network data representation. Traversal time for each link in each layered network is computed by a link traversal function, which is time-dependent.

The Activity Generator provides information on mode preferences for each trip. Conceptually, this information is captured in simple, alphabetical expressions. For example, a trip that contains a walking leg from a traveler's house to his/her car, a car leg to parking at the place of work, and a walking leg from the parking lot to his/her actual work location would be represented as *"w-c-w."* The Route Planner uses this information in the following way: For the first leg of the trip, the walking leg, the Route Planner searches for possible paths within the walking layer of the network to obtain a walking route from the home to the parking location of the individual's vehicle. As soon as the walking path is found, a series of least-cost driving links is found to obtain a route to a parking location near the work location. A walk route is then developed to move the traveler from the parking lot to the work activity location.

The last two legs of the above route highlight the capabilities of the Route Planner. Once the search algorithm is in the car network layer, it chooses additional links from the car network or parks the vehicle and chooses links from the walking network—whichever is lower in cost. The Route Planner ensures that the final link is a walking link in this example. Trips that cannot be feasibly planned are marked and provided as output from the Route Planner to be fed back to the activity list to choose a new activity time or location or mode for travel to the activity.

3.4 Traffic Microsimulator Module

The Traffic Microsimulator module executes individuals' travel plans link by link, as provided by the Route Planner, at the start time specified by the plan. Plans that overlap in time are executed simultaneously by the Traffic Microsimulator, resulting in overall transportation dynamics along the intermodal networks.

3.4.1 Traffic Microsimulator Major Input/Output

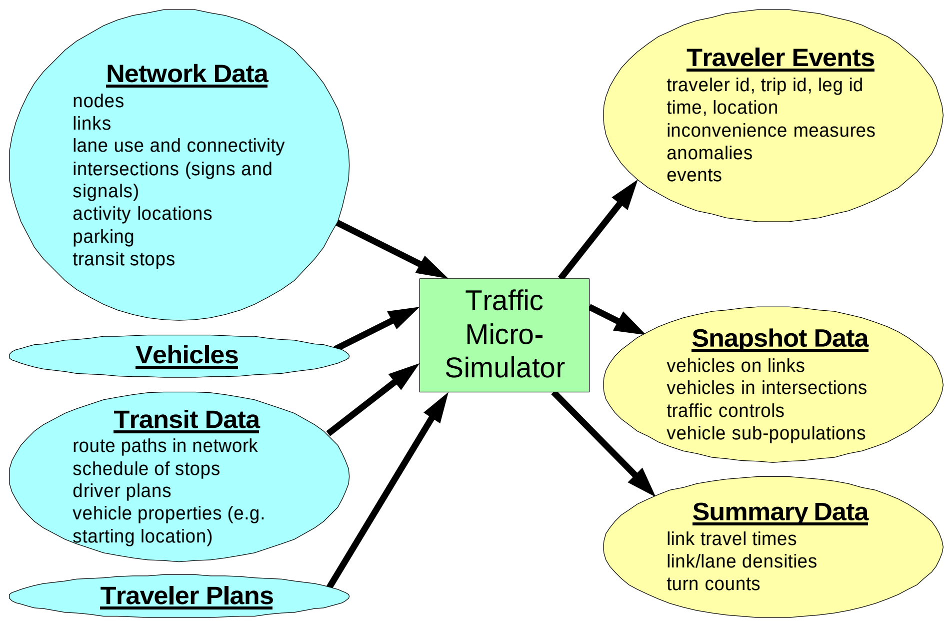

Figure 11: The major inputs to the Traffic Microsimulator are the intermodal plans, transit data, TRANSIMS network data, and vehicle data.

The Traffic Microsimulator uses the TRANSIMS network and the intermodal plans as diagrammed in Figure 11.

At a minimum, TRANSIMS network information must include the locations of streets and intersections, the number of lanes on the streets, the manner in which the lanes are connected, and some parking locations on the streets at a collection of activity locations. Some studies might benefit from or require more detailed information about the network. Examples of more detailed street information that the Traffic Microsimulator is capable of using include the presence of turn pockets and merge lanes, lane use restrictions (such as high-occupancy vehicle lanes), turn prohibitions, and speed limits. Each intersection has a controller associated with it, which can vary from a stop or yield sign to a traffic signal, or to a set of coordinated traffic signals. Another kind of network information is the list of transit stops serviced by each route. The actual transit schedule is encoded in the travel plans of transit drivers. Transit drivers stop to pick up or drop off passengers at transit stops.

For each traveler, the simulation must have a complete description of the traveler's transportation plans. A plan is broken down into a sequential set of trips, which must begin and end at an activity location, such as home, work, or shopping center. A trip is further decomposed into a set of unimodal legs. A traveler can use only a single mode of transportation on a leg. Accordingly, several legs are chained together to form a single trip.

3.4.2 Traffic Microsimulator Description

As a simulated day progresses, each person moves from one activity to the next according to the predefined plan, using combinations of modes such as walking, driving a vehicle, and/or riding in a (private or public) vehicle. The Route Planner provides the link-by-link travel plan of the traveler including the mode of travel. All TRANSIMS vehicles are simulated in detail to include driving on roads, stopping for signals, accelerating, decelerating, changing lanes, stopping to pick up passengers, and so on. Mode changes—for example, from walking to car or to transit—are explicitly simulated based on information contained in the traveler's plan.

Vehicles follow a simple set of rules that guarantee that no collisions will occur. Phenomena such as reaction times and limited visibility are not simulated explicitly. However, the effects of these phenomena are simulated by the values of parameters used in the driving rules so that the fundamental flow-density diagram matches real, observed traffic.

The simulation can estimate the impact of hypothetical changes on quality of service. It provides answers to such questions as the following:

- What would be the effect on traffic patterns of building a proposed highway?
- What is the impact of a change in transit schedules on riders?
- Can changing signalization alleviate congestion?
- Are there common demographic characteristics of the subpopulation most affected by a particular infrastructure change?

The analytical power of the Traffic Microsimulator lies in aggregating the results of millions of interactions within the transportation system. These summaries focus on issues of interest such as congestion and many other issues that can be addressed based on understanding very detailed traffic flow information.

3.4.2.1 Single Trip Example

This example considers a work-to-home trip. It begins and ends at activity locations that are coded in the TRANSIMS networks.

- Leg 1: walk from activity location *W* to bus stop *X*
  where
  W is the designation of the work activity location, and
  X is the designation of a bus stop in the network description
- Leg 2: take route *Y* to bus stop *Z*
- Leg 3: walk to parking lot *P*
- Leg 4: drive to day care at activity location *D*
- Leg 5: drive, with one passenger, to parking location *P2*
- Leg 6: walk to activity location *H* (home)

The first leg of this hypothetical trip is a *walking leg*. For walking legs, TRANSIMS does not explicitly microsimulate the second-to-second locations of pedestrians. Based on information contained in a given traveler's plan, legs of a trip that involve walking will begin at the time specified by the Route Planner at the point of origin for the leg and end at the destination of the leg. The traveler arrives at the destination at a simulation time computed by adding the delay time (contained in the plan) to the start time for the walk mode leg. No additional information is required or generated for walk mode legs.

The second leg of the trip is a *bus leg*. A bus leg plan requires only one additional piece of information: the acceptable route(s). The precise itinerary of the bus the traveler gets on is determined by the driver's plan. The traveler simply waits until his desired stop is reached and exits the bus. Bus loading and unloading is represented explicitly by the microsimulation. Resource constraints such as vehicle capacity and parking lot capacity are observed. If a bus is full when it reaches the bus stop, the traveler is not permitted to board and will wait for the next bus on the same route. With this level of detail, it is easy to determine how many passengers cannot find space on the bus, or how many minutes a traveler must wait for a bus.

The next leg calls for the traveler to walk to the parking lot after getting off the bus. In this instance, the parking lot is where the traveler left his/her private vehicle. This walking leg is handled as previously described.

Upon arriving at the parking lot, the traveler is associated with a specific vehicle, which must either have been left in the parking lot earlier in the simulation or placed there during initialization. The traveler and car exit the parking lot into the traffic network. The traveler's plan specifies exactly which turns he/she will take until he/she arrives at the daycare center. Here, the traveler waits until his/her passenger enters the car. The passenger's plan will specify what vehicle to ride in, and the passenger will be waiting for this vehicle to arrive. Then, as before, the driver re-enters the transportation network, following his/her plan, parks, and walks home.

As described above, a transit schedule is implemented by providing plans for each transit driver. Like other travelers, a driver can switch vehicles, switch routes (with or without switching vehicles), and take layovers of prescribed duration or ending time at specific control points.

The Traffic Microsimulator enforces physical constraints—travelers cannot be in two places at once, and they cannot create vehicles. Thus, travelers and vehicles will be initialized and placed in their initial start locations based on information in the plan file.

3.4.2.2 Cellular Automata

Vehicle movement is accomplished using a cellular automata simulation. Each section of roadway (see Figure 12) is divided into cells of a grid 1-lane wide by 7.5-meters long. Each cell either contains a vehicle (or a part of one) or is empty. The simulation is carried out in discrete timesteps, each simulating one second of real time. On each timestep, a vehicle on the network decides whether to accelerate, brake, or change lanes in response to the occupancy of nearby grid cells. After every vehicle is allowed to make these decisions, they are all moved to new grid cells in accordance with their current velocity.

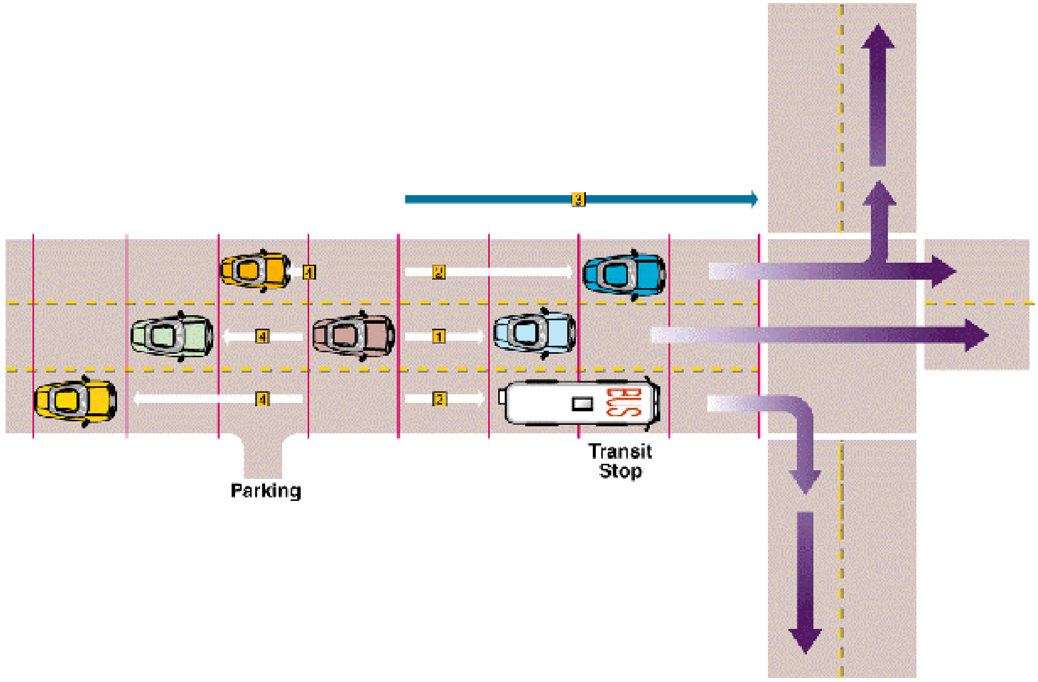

Figure 12: The roadway is divided into lanes, and the lanes are divided into cells. Each lane's possible turning movements at intersections are defined. Controls at intersections manage traffic and maintain information about oncoming vehicles. Shown from the perspective of the middle peach-colored car, the gap (labeled 1 in the figure) between vehicles determines accelerations and influences lane changes. The gaps (2, 4 for left-lane change, 6 for right-lane change) to vehicles in other lanes also influence lane changing. The distance at the intersection (3) determines the relative importance of changing lanes. The gaps (4, 5, 6) to upstream vehicles from a parking facility determine whether vehicles can exit the parking facility.

As part of the lane-changing procedure, transit vehicles scan the nearby cells for transit stops, which they must service. The transit vehicles can examine the queue at the stop to see if anyone is waiting for the vehicle and can also query their passengers to see if anyone wants to get out at the stop. If so, they either stop in the cell next to the transit stop or pull off the grid, depending on the type of transit stop. Passengers take a fixed time to enter or leave the vehicle.

Each intersection has traffic control logic, which controls entry of vehicles into the intersection. Traffic controllers examine the traffic in each lane at the intersection. If the intersection is clear, vehicles pass through it in a fixed amount of time and are placed on the next roadway's grid.

Vehicles entering the roadway from parking locations or off-street transit stops can enter any lane with a large enough gap between it and the oncoming traffic. The gap must be large enough to ensure that on the next timestep, no vehicles will collide with vehicles entering the roadway.

The simulation guarantees that each vehicle makes decisions based on the state of every other vehicle in its local vicinity (i.e., 5 cells) at the same time. That is, every vehicle in the network makes its acceleration decision based only on information available at time *t*, which does not include the time *t+1* positions of vehicles that have already made their acceleration decision. This parallel update scheme ensures that the simulation results do not depend on the order in which streets in the network are updated.

To accomplish the simulation update, a single timestep must be broken down into several parts. First, the vehicles at the head of a queue waiting to leave a parking location are allowed to enter the road. Then, all vehicles are allowed to change lanes (in this step, transit vehicles are allowed to enter transit stops also). In order to avoid possible collisions when vehicles in two lanes at time t both try to change into the same lane, we alternate the direction of lane changing every timestep. Next, those vehicles that will enter intersections are marked and given instructions from the intersection controller about the availability of their destination link. Every vehicle makes its acceleration decision, and all of the vehicles are moved. Finally, vehicles are allowed to exit into parking locations.

3.4.2.3 Traffic Microsimulator Output

There are four major kinds of output from the Traffic Microsimulator. Figure 13 shows summary data (both spatial and temporal), traveler events, and snapshot data. Spatial summaries include data aggregated over user-defined sections of roadway defined along the street networks—for example, densities and total flow in a 150-meter section. Temporal summaries include data about travel times along streets at various times of day. Almost anything that happens to a traveler can be reported as a time-stamped event in the event output. Commonly used events include begin/end waiting at a given location, such as a bus stop, begin/end a leg, pass through an intersection, and enter a vehicle. If desired, traffic animation can be produced from the snapshot files, which contain time, position, and velocity information for each vehicle in the simulation. These files can also be used to recover data that has not already been provided in the summary data files. For instance, if some new study requires the average gap between vehicles, it can be computed from snapshot data.

Dumping out the snapshot data for a 24-hour simulation of a major metropolitan area creates extremely large files. Users are allowed to restrict output to smaller portions of the network and specific times during the simulation, as well as to select only a few desired fields or only those records that meet certain criteria. For example, a user may choose only specific events like beginning a leg, only particular travelers, or only vehicles traveling above a given speed. The sampling rate and reporting frequency for each of these data types is controlled by parameters selected by the user.

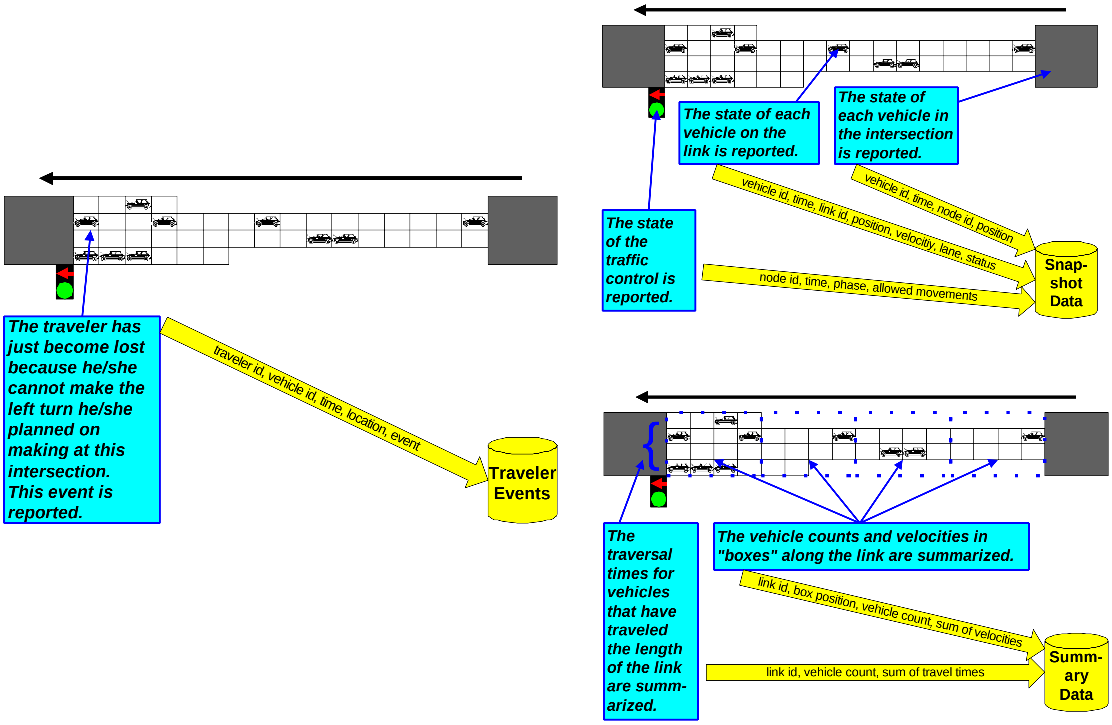

Figure 13: The various types of Traffic Microsimulator output—traveler events, snapshot data, and summary data—are depicted.

3.5 Emissions Estimator Module

The Emissions Estimator module translates traveler (vehicle) behavior into estimates of air quality, energy consumption, and carbon dioxide emissions. Four major computational submodules compose the Emissions Estimator module: (1) emissions, (2) atmospheric conditions, (3) local transport and dispersion, and (4) chemical reactions. The MODELS-3 system, developed by the EPA, contains the last three of these and is not part of the TRANSIMS framework.

3.5.1 Emissions Estimator Major Input/Output

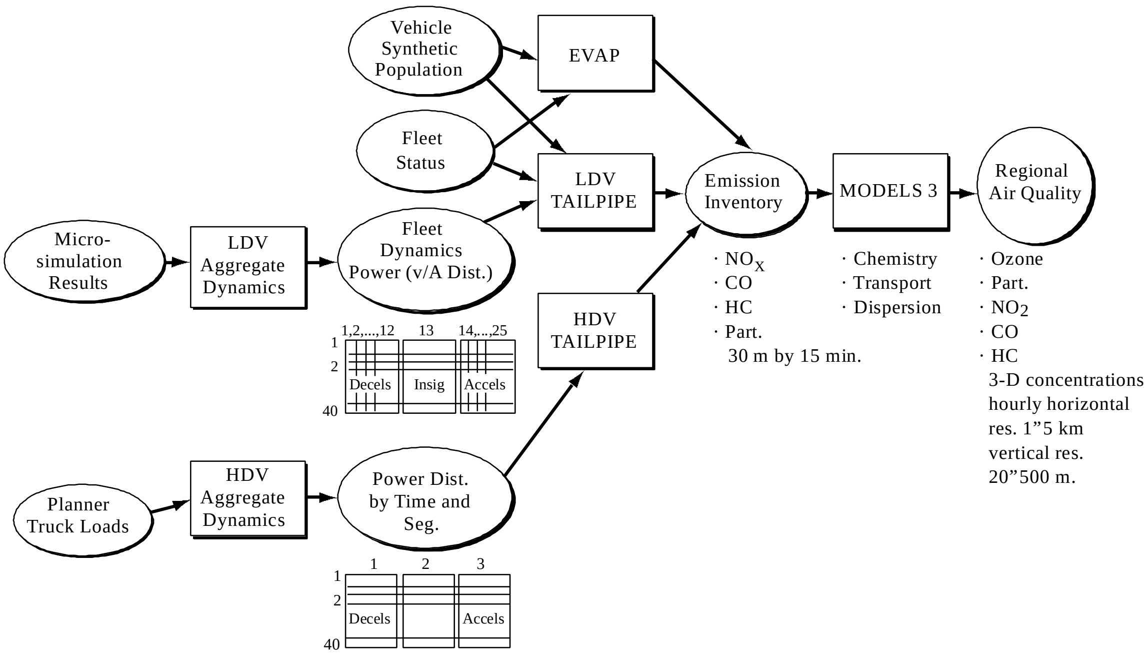

Figure 14: Emissions Estimator data flow.

The data flow of the Emissions Estimator is summarized in Figure 14. The Emissions Estimator requires information on the fleet composition developed from the Population Synthesizer, vehicle loads, and traffic patterns. The Population Synthesizer provides vehicle fleet characteristics including the fraction of the fleet that is malfunctioning. The traffic patterns are produced by the Traffic Microsimulator.

The bulk of the information required by the Emissions Estimator is produced by the Traffic Microsimulator. The following items are included in this category:

1) spatial summaries of vehicle velocities over 30-meter sections of roadway

2) summaries of the gaps between vehicles as a function of vehicle speeds

3) histograms of the number of vehicles entering a link grouped by velocity-acceleration product summed over time since the vehicles were parked

4) bus loads

Emissions Estimator output data is aggregated on 30-meter segments for each simulated 15-minute period. Fuel economy and CO_2 emissions are also estimated. The emission inventory is designed to be used with the MODELS-3 code, developed by the Environmental Protection Agency (EPA) to produce three-dimensional, hourly, gridded emissions values over the metropolitan area.

3.5.2 Emissions Estimator Description

There are three major submodules in the Emissions Estimator:

1) Fuel Evaporation submodule

2) Light-Duty Vehicle (LDV) Tailpipe submodule

3) Heavy-Duty Vehicle (HDV) Tailpipe submodule

The Fuel Evaporation submodule treats emissions associated with resting losses, running losses, hot soaks, and diurnal pressure changes. It deals with both normally operating vehicles and with vehicles with significant leaks in the fuel system.

The LDV Tailpipe submodule treats tailpipe emissions from cars, small trucks, and sport utility vehicles. Important aspects include:

1) malfunctioning vehicles

2) emissions from cold starts

3) emissions from warm starts in which the engine is still warm but the catalyst is cold

4) emissions from off-cycle conditions, conditions that occur outside those in the federal test procedure, or that render the pollution controls inefficient

5) normal driving

The HDV Tailpipe submodule treats tailpipe emissions from trucks and buses. While truck emissions are not sensitive to power levels as are LDVs, their emissions are sensitive to the load carried by the vehicle.

3.5.2.1 Fuel Evaporation Submodule

The Fuel Evaporation submodule uses vehicle activity information generated by TRANSIMS and activity-specific emissions models from the EPA to generate the evaporative emissions values. It uses an age distribution of the vehicles in the simulation and includes the ratio of light-duty trucks to automobiles. The evaporative model provides values for outputs of hydrocarbons over 15-minute periods (interpolated from hourly values) on 30-meter segments along links, or at parking locations.

The Fuel Evaporation submodule uses information from the Traffic Microsimulator to determine the location of a vehicle and whether it is presently operating or has operated in the previous hour. If the vehicle is not operating or has not operated in the last hour, resting losses and diurnal evaporative emissions are calculated for the vehicle using the vehicle categories and the equations developed for the Mobile 6 model of the EPA. These emissions are then assigned to the spatial grid where the vehicle is located. If the vehicle is operating, running losses are calculated using the Mobile 5 formulation. If a vehicle has operated within the past hour, hot soak emissions are calculated using the hot soak model developed by the EPA for Mobile 6.

For vehicles in the simulation area but not operating during the simulation period, diurnal and resting losses are calculated and assigned to the appropriate spatial location.

3.5.2.2 Light-Duty Vehicle Tailpipe Submodule

The Light-Duty Vehicle Tailpipe submodule treats emissions from off-cycle driving, malfunctioning vehicles, normal driving, idling, and vehicles with cold engines and/or catalysts. The LDV data flow is illustrated in Figure 15. Three major sets of information—fleet composition, fleet status, and fleet dynamics—are inputs to this submodule.

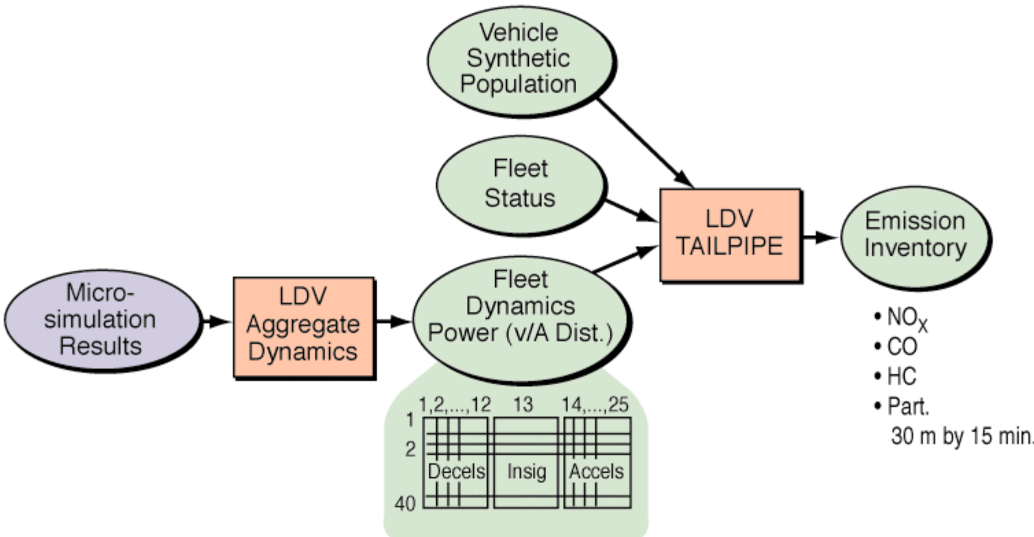

Figure 15: LDV Tailpipe submodule data flow.

3.5.2.2.1 LDV Tailpipe Fleet Composition

Fleet composition is given in the vehicle file from the Population Synthesizer. The 23 categories of vehicles in this file reflect such characteristics as engine-to-weight ratio and the catalyst type.

3.5.2.2.2 LDV Tailpipe Fleet Status

Fleet status is developed from the pattern of usage of the vehicles traversing a given link. Fuel consumption is a function of the temperature of the engine and the catalyst. The Traffic Microsimulator records when and where the vehicles have been operating on an individual-vehicle basis. Consequently, the fraction of the fleet in various phases of cold engine/cold catalyst operation can be estimated from the vehicle origins and links traveled.

3.5.2.2.3 LDV Tailpipe Fleet Dynamics

Given good estimates of a vehicle's continuous power output, its velocity, and the type of vehicle (provided in the synthetic population data file), predictive models of tailpipe emissions can be developed. The principal component in estimating a vehicle's continuous power is directly proportional to the product of its velocity and acceleration. Thus, good estimates of these variables for the vehicles in a TRANSIMS microsimulation are extremely important to predict emissions.

The velocity of TRANSIMS LDVs is estimated in the following way: Since the microsimulation is a cellular automata, only a discrete number of vehicle velocities are possible. The microsimulation vehicles can be moving an integer number, 0 to 5, cells per second. These represent speeds of 0 to $7.5*5=37.5$ meters per second (or 0 to 84 miles per hour). For the Emissions Estimator, we summarize TRANSIMS velocity output over 15-minute intervals and 30-meter segments (boxes) by counting the number of vehicles in each of the six speed bins, 0 to 5 cells per second.

Because of the significant nonlinear character of the relationship between power and emissions, a more finely grained estimate of velocity is needed. This is accomplished by fitting a continuous model to the discrete data in the six speed cells that range from 0 to 37.5 meters per second. From this fit, the number of vehicles traveling at speeds in 2-mph increments is determined.

Next, vehicle accelerations are estimated. Data collected in the EPA's *three-city* study, where many vehicles were fitted with a data logger that recorded times and speeds throughout the vehicles' travels, are used. Frequency distributions of accelerations for given velocities were determined from these data. From these data, it was evident that the fraction of vehicles having a velocity-acceleration product greater than a specified level falls off exponentially with the level.

In the case of decelerations, the frequency also falls off exponentially with the velocity-deceleration product. These relationships form one of the principal empirical underpinnings in the approach used here. All accelerations are placed in one of three groups: (1) hard accelerations, (2) insignificant accelerations, and (3) hard decelerations. From empirical data and analysis, hard accelerations are defined as those accelerations greater than the accelerations associated with the 10% of vehicles that have the largest velocity-acceleration product. The number of vehicles undergoing a hard acceleration is estimated, and 12 different power levels above the hard acceleration cutoff are chosen to represent different levels of aggressiveness. The total population of vehicles undergoing hard acceleration from a given speed is then distributed over the 12 power (or equivalently acceleration) levels. This data is used to estimate the relative proportions of vehicles having different levels of acceleration within the group of vehicles that undergo hard accelerations.

A second methodology to estimate the number of vehicles undergoing hard-acceleration, insignificant acceleration, and hard deceleration is being investigated. It is hypothesized that the fraction of vehicles undergoing hard acceleration would be related to average speed changes between segments or to the standard deviation of speeds within a segment. If most vehicles are accelerating so that the average speed is increasing along the link, a higher proportion of hard accelerations would be expected. On the other hand, if the average speeds are not changing but there is a large standard deviation of speeds within the segment, slower vehicles would be accelerating to regain their desired speed after encountering a temporary traffic jam. This hypothesis was tested by examining data collected by the California Air Resources Board. Here, driving on freeways with seven different levels of congestion and driving on arterials with three different levels of congestion were described. A simple two-parameter curve-fit was used to relate the probability of a hard acceleration to gradients in speed to the standard deviation of speeds within a link segment. In the calibration, only the fastest freeway, a mid-speed freeway that had the highest standard deviation of speeds, and the fastest arterial that had a large gradient in speeds for vehicles leaving a signalized intersection were used. The relationships on the remainder of the arterials and freeways were tested, and good results were obtained. The relationships are used to estimate the fraction of the vehicles undergoing hard acceleration, insignificant acceleration, and hard deceleration. The fraction of vehicles in the hard-acceleration group is distributed with the exponential distribution into 12 groups of velocity-acceleration product. Emissions are estimated for each velocity grouping (2-mph bins) and each acceleration grouping (25 levels of 1.5 feet per second).

3.5.2.2.4 LDV Emissions Model

Barth et. al. [4] developed an improved modal emission model for LDVs. They carried out extensive tests on over 300 vehicles chosen to represent the major types of emitters in the existing LDV fleet. They also worked with other data to help draw associations between the tested vehicles and the fleet at large.

Using speeds and accelerations, the model computes the tractive power by taking into account engine friction losses, rolling resistance, wind resistance, changes in kinetic energy, and changes in potential energy. It also considers the power to drive accessories such as air conditioning, and it estimates drive-train efficiency. With the engine power known, it calculates the rate of fuel consumption and engine emissions. It treats enrichment, enleanment, and stoichiometric operations, as well as cold-start operation.

Once the engine emissions are calculated, catalyst pass fractions are used to calculate the tailpipe emissions. The approach uses a composite vehicle to represent vehicles in the same class. A regression approach was used to define the parameters required by the model. The vehicles were all tested over cycles involving very high power demands and a variety of driving patterns.

There are composite vehicles representing normal-emitting cars categorized by technology, low and high power-to-weight ratios, and mileages above or below 50,000. The technology categories are (1) no catalyst, (2) two-way catalyst, (3) three-way catalyst with carburetion, (4) three-way catalyst with fuel injection, and (5) Tier 1. Only the last two categories are broken into mileage or power-to-weight ratio groupings. There are high-emitting composite vehicles for technologies (3) through (5), but they are not further subdivided into power-to-weight ratios or mileage groupings.

There are composite vehicles representing normal-emitting trucks with model year categories: (1) pre-1979, (2) 1979 to 1983, (3) 1984 to 1987, (4) 1988 to 1993, and (5) 1994 and newer. In age categories (1) through (3), there is only a single composite vehicle. For age category (4), there are categories for trucks above and below 3750 pounds loaded-vehicle weight, while for category (5) there is a category for trucks with loaded-vehicle weights between 3751 and 5750 pounds, and a category for gross vehicle weights between 6001 and 8500 pounds. There are composite vehicles representing high-emitting trucks for model years 1984 to 1987, 1988 to 1993, and 1994 and newer. In the high-emitting category, there are no breakdowns by vehicle weight.

3.5.2.3 Heavy-Duty Vehicle Tailpipe Submodule

The Heavy-Duty Vehicle (HDV) Tailpipe submodule treats emissions from trucks and buses as they idle or deliver loads about the city. Truck emissions are sensitive to power demands associated with accelerations, climbing grades, and hauling heavy loads. Bus emissions are sensitive to power demands, but loads play a relatively smaller role. Figure 16 describes the data flow for the HDV Tailpipe submodule. There are three major sets of information that must be developed: (1) fleet composition, (2) fleet status, and (3) fleet dynamics.

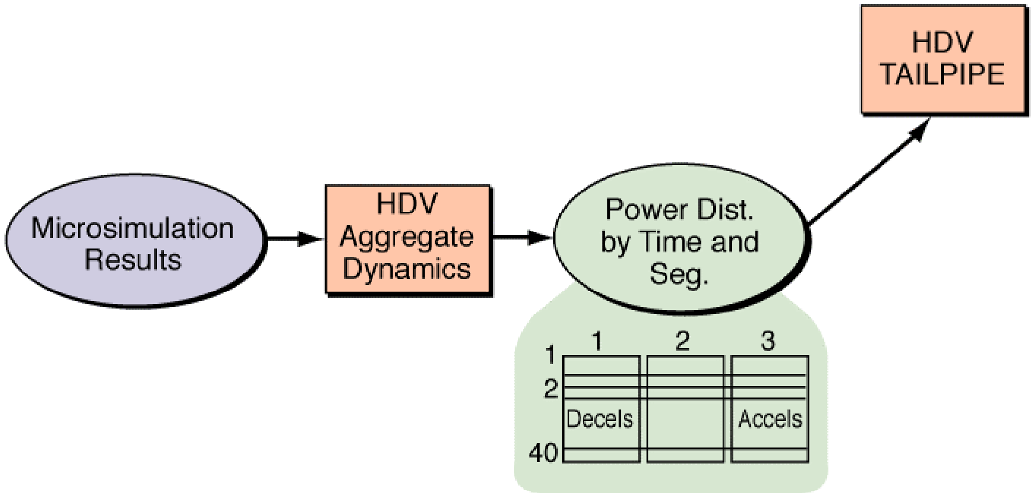

Figure 16: HDV Tailpipe submodule data flow.

3.5.2.3.1 HDV Tailpipe Fleet Composition

The HDV Tailpipe submodule needs the number of buses and trucks. For both trucks and buses, the fractions in various categories that describe engine size, chassis size, and model year are needed. In the case of buses, if the transit authority has purchased buses with better emissions controls than required for the corresponding model year, the characteristics of the additional controls are also required.

3.5.2.3.2 HDV Tailpipe Fleet Status

The HDV Tailpipe submodule requires an estimate of the vehicle loads as the HDV travels about the city. Currently, there is not a malfunctioning vehicle class, nor is there a cold-engine class.

3.5.2.3.3 HDV Tailpipe Fleet Dynamics

Buses and trucks have very low accelerations and are usually driven at full throttle whenever the speed is less than desired and there is adequate headway to accelerate. Consequently, the HDV aggregate dynamics submodule produces only three levels of acceleration: (1) maximum deceleration, (2) constant speed, and (3) maximum acceleration. The maximum acceleration is a function of the engine size, the grade, and the total vehicle weight.

The Traffic Microsimulator provides populations of each type vehicle in each speed bin for each 30-meter segment of each link over each 15-minute period of the simulation. It also provides the fraction of the time an HDV has a gap in front of it and the fraction of the HDVs that have plans requiring a stop in any given 30-meter segment. The HDV aggregate dynamics submodule estimates the fraction of vehicles in each speed bin that needs to accelerate and have the opportunity to accelerate and assigns them the maximum acceleration for each type and load of vehicle. It also estimates the number and type of vehicles that will not accelerate and the number of vehicles that must decelerate in a segment in a speed bin.

3.5.2.3.4 HDV Emissions Model

West Virginia University investigators have carried out detailed studies using a chassis dynamometer on a variety of HDVs. For the TRANSIMS work, they have developed emissions values for two different classes of buses under full throttle, constant speed, and deceleration in 2-mph increments from idle to 78 mph. They have also developed full-throttle trajectories for the buses and have provided relationships to correct the trajectories for vehicle weight and grade.

The HDV model estimates emissions for the HDVs with specified loads, speeds, grades, and acceleration categories. The principal emissions of interest are NO_x and fine particles, and these are provided as 15-minute averages for each 30-meter segment of each link.

3.6 Selector Component

The process of iterative feedback is an important feature of TRANSIMS. It is the primary mechanism used to achieve internal consistency (i.e., to achieve a reasonable agreement among the travel demands expressed in the activities list, the travel plans to meet these demands, and the execution of the plans in the Traffic Microsimulator) among the various computational modules. It selectively feeds back information from the Traffic Microsimulator into the Activity Generator and Route Planner or from the Route Planner to the Activity Generator. In effect, this information is used to modify a designated subset of activities and/or plans of the synthetic population to achieve realistic overall traffic results. In addition, it allows the overall computational system to abstractly reflect *learned* behavior within the simulated population represented. TRANSIMS synthetic travelers do not *learn* in the way humans do, but this technique emulates the ability of humans to learn from day-to-day experiences in order to avoid congestion, etc. It also provides a way to model intelligent responses to information that may be provided ultimately by Intelligent Transportation Systems (ITS) technologies.

The Selector controls the iterative process. A typical TRANSIMS study involves repeated iteration between components such as the Activity Generator, Route Planner, and Traffic Microsimulator. There is no single, *standard* selector component, however, because different study designs involve different iteration schemes. A variety of selectors have uses in different studies or other contexts.

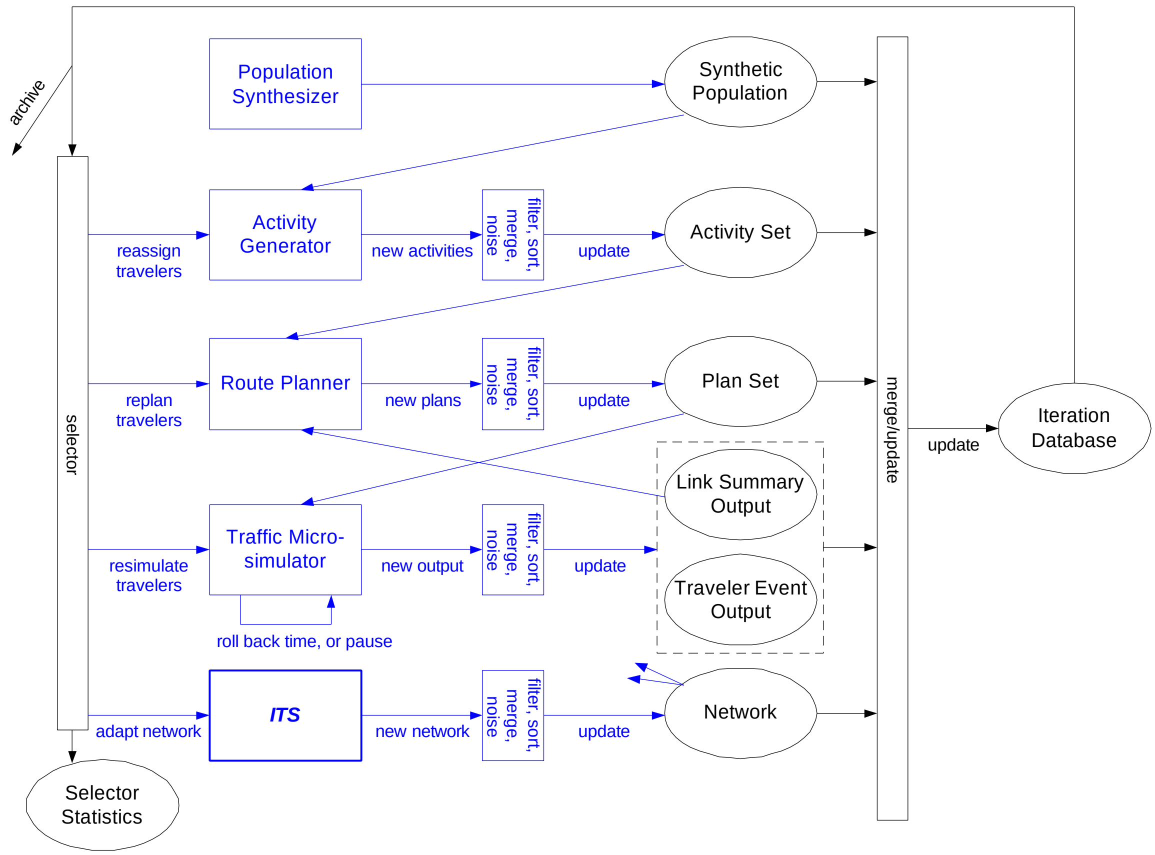

Figure 17: TRANSIMS framework diagram with the Selector at left.

The user can prepare a script to control the whole process of iteration. The script uses a special control language specifically developed for this component. It allows the user to filter results, run repeated iterations, establish stopping criteria, and perform a variety of other operations that make the analyst's job less manpower-intensive.

At the beginning of each iteration, the iteration script controlling the current study typically invokes the Selector. (The script might even use a different selector for each iteration in a study.) When the Selector runs, it usually will do the following:

- Read information about the travelers from the iteration database.
- Examine each traveler and decide whether to
  - regenerate his/her activities using the Activity Generator,
  - choose a new route between his/her existing activities using the Route Planner, or
  - retain his/her existing activities and the planned route between them.
- Write the selections made for each traveler into data files that can be read by the Activity Generator and/or Route Planner when they are executed.
- Summarize the selections made and the current state of the system into the Selector statistics data file.

Figure 18 illustrates the Selector's decision-making process.

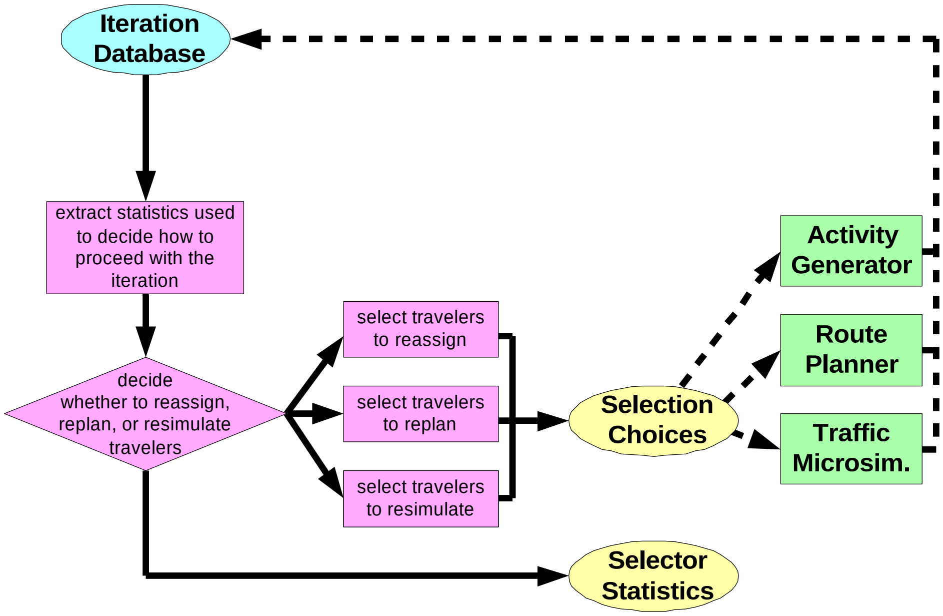

Figure 18: The Selector's decision-making process.

After the Selector completes the selection process for all of the travelers, the Activity Generator, Route Planner, or Traffic Microsimulator runs to calculate the updated activity set, plan set, or microsimulation output files, respectively, according to the decisions made by the Selector. The iteration script will reinvoke the Selector again at the start of the next iteration in the study.

The flexibility of the TRANSIMS framework allows for countless variations on the selection process described above. For example, in some studies, part of the Selector may run again *after* the Activity Generator or Route Planner completes its execution in order for the Selector to decide which of the activities or plans just generated will be accepted for travelers. Those not accepted are discarded, and the previous activities or plan is retained. One can also design selectors that will feed travelers to the Activity Generator or Route Planner one-by-one so that the Selector, Activity Generator, Route Planner, and Traffic Microsimulator all execute simultaneously with their coordination controlled by the Selector. This may increase the computational efficiency of a study and allow for new experimental designs with finely controlled iteration. The Selector might even make additional choices such as

- which version of the Activity Generator, Route Planner, or Traffic Microsimulator will run during the present iteration
- whether transit schedules will be adjusted or vehicles added or removed from the transit fleet
- whether network characteristics like traffic signal timing, congestion pricing, or roadway information signs will be altered
- which travelers receive data from traffic information systems
- whether to complete the study (i.e., end the iteration) because the iterations have converged sufficiently (or diverged)

Several Selector implementations have been written that have use in typical transportation planning studies. For example, Figure 19 shows a typical iteration scheme that is set up by the Selector script. In this scheme, activities, plans, and microsimulations are iterated until traffic behavior on the network stabilizes. It is not difficult for analysts to write additional selectors for their own specialized studies.

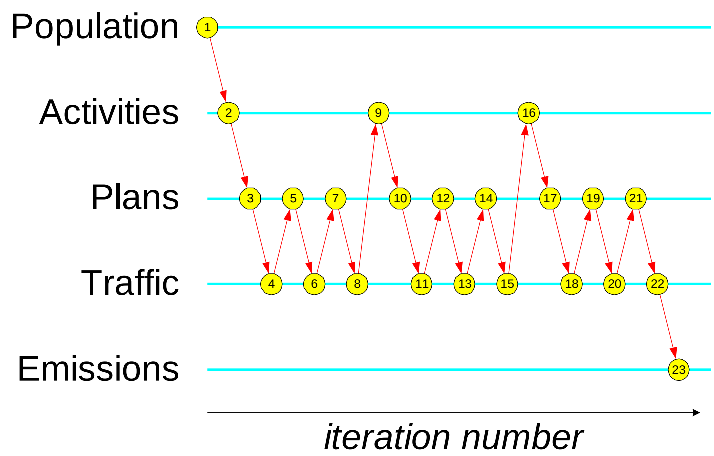

Figure 19: A typical iteration scheme set up by selector scripts.

The major input to the Selector is the iteration database. It contains a summary history of each traveler's attributes, expectations, and experiences during the iterations within a study. The Selector uses these data items to make its selection decisions. *Attributes* represent information about travelers like their age, income, gender, or profession. *Expectations* encompass information such as how long a traveler expects to travel between two activities based on the route between them generated by the Route Planner. *Experiences* compose information extracted from detailed Traffic Microsimulator output—for instance, the actual travel time realized in the microsimulation between two activities. The analyst may choose which attribute, expectation, and experience data reside in the iteration database for a particular study. These data form the universe of information readily available to the Selector; however, additional data from activity sets, plan sets, and microsimulation output might also be used by some selector implementations.

The Selector also has two principal outputs: selector statistics and selection choices. The *selection choices* files simply list the travelers that will be reassigned activities, replanned, resimulated, etc. These files embody the detailed decisions of the Selector. The *selector statistics* provide a basic summary of the choices the Selector makes; e.g., how many travelers are being replanned, distributions of the difference between expected travel times and experienced travel times for various traveler populations, and the like.

4. ANALYSIS: NEW ANALYTICAL TECHNIQUES

Most traditional analyses can be performed by using the aggregate, summary data that result from the TRANSIMS microsimulation. In an existing case study for Dallas (available at *http://transims.tsasa.lanl.gov*, for example, we used the Traffic Microsimulator output to compare two alternative improvements to the highway infrastructure in the vicinity of two major malls. Ostensibly, these improvements were intended to alleviate the congested conditions along a major corridor adjacent to the malls. The types of questions addressed were "What is the average travel time for all travelers going to the mall for each of the three infrastructures considered?" or "To what extent do mean freeway speeds increase under each of the choices considered?" But TRANSIMS also affords users the ability to provide new and different insights into the more fundamental, dynamic nature of the transportation system. The TRANSIMS project has developed and successfully demonstrated feasible technical approaches for conducting the following analyses.

4.1 Variability Analyses

Aggregate methods normally depend on averages and expected values of such key variables as link travel times and link speeds. These averages pertain to all vehicles using the links during a given time period. Estimates for the averages are based on simple models that relate traffic volume to the average speed of all vehicles using a particular link. In contrast, TRANSIMS uses local interactions of simulated vehicles (updated every second) to *generate* instantaneous vehicle speeds that can be averaged over user-defined time periods. These averages are sensitive to signal timing, local traffic conditions, ITI options, etc. Because travelers are tracked throughout, TRANSIMS can produce variability measures by comparing the impact on individual travelers or subpopulations of travelers with overall population trends. Statistical measures generated by these methods can address questions such as the following:

- What is the variance, or spread, in total travel times for the subpopulation of mall travelers between the infrastructure changes being considered?
- Given that the average travel time for the subpopulation of mall travelers improves, does the variance or spread of their travel times improve also, and if so, by how much?

Variability measures are extremely important. They may be used (1) to assess infrastructure reliability, (2) in equity analysis, and (3) to assess uncertainty induced by the model components. These variability measures are discussed in the following sections.

4.1.1 Infrastructure Reliability Analyses

The basic notion is to define the network reliability in terms of its dependability in meeting traveler expectations from day to day. For example, if 5% of the population experiences a dramatic variation in the time it takes to make the same trip, using the same route, under the same conditions from one day to the next, one could subjectively conclude that these travelers might have a high degree of dissatisfaction. These travelers are probably not very confident that they can accurately predict, from day to day, how long their trips will take. This measure of unpredictably, which is easily computed by TRANSIMS, is relevant and manifests itself in real traffic systems—especially those at or near maximum capacity. Having the capability to compare proposed infrastructure alternatives using measures such as reliability is unique to TRANSIMS and other simulations that track individual travelers. These measures will allow planners to focus better on that subset of the population likely to be dissatisfied because of an inability to make a reasonable prediction of the time it takes to make a trip. If the percentage of travelers in this category is large enough and the variability in travel time is big enough, these factors are likely to be of greater importance to planners than improving the average travel time for the overall population by a marginal amount.

4.1.2 Equity Analyses

Equity analysis is an important aspect of any travel network study. The subpopulations used for equity comparisons may be based on travel characteristics such as origins or destinations, length of trips, or traveler demographics. The availability of tracking information on individual travelers in TRANSIMS makes partitions of the populations based on any of these factors very simple to create. Because TRANSIMS includes environmental and other kinds of impact analyses, one also can begin to measure effects such as the total pollution or the energy utilization attributable to a given subpopulation such as the mall travelers. While traditional equity analyses usually discriminate subpopulations based on information intrinsic to the supporting databases (e.g., zone of origin, type of trip, etc.), TRANSIMS provides an interesting capability to discriminate subpopulations based on dynamic information resulting from the microsimulation. For example, one could ask if there is any correlation between the subpopulation of travelers experiencing a large variability in travel times from day to day, where they live, what routes they have in common, etc.

4.1.3 Uncertainty Analyses

Like all simulations or models, TRANSIMS only approximates reality. Accordingly, it is imperative for analysts and decision makers to understand how much of the change in simulation results is attributable to the structural properties of the computational components of the simulation. These types of effects are defined as *uncertainty* in simulation results induced by the computational system itself. For example, if a particular algorithm represents a random process, one would like to know how changing a random number *seed* would affect the simulation results. If these effects are large compared to effects attributable to an infrastructure change, then, on the basis of the simulation results, one should be much less confident that the infrastructure change will make any difference. There has been a significant effort in the first case study to define and quantify the effects of uncertainty.

4.2 Output Visualizer Module

One of the most important analysis techniques is graphical visualization of the results of a study. The TRANSIMS Output Visualizer module allows the analyst to view dynamically the output from the Activity Generator, Route Planner, and Traffic Microsimulator modules. All displays are both temporally and spatially dynamic. That is, spatial areas may be located and the display *zoomed* to give better viewing. Additionally, all displays may be stepped through time to assess the changing characteristics of the roadway network. Figure 20 shows a view from the tool of a single route from the Route Planner. Planned routes are overlaid in time and displayed in Figure 21. These show potential congestion points on the roadway network.

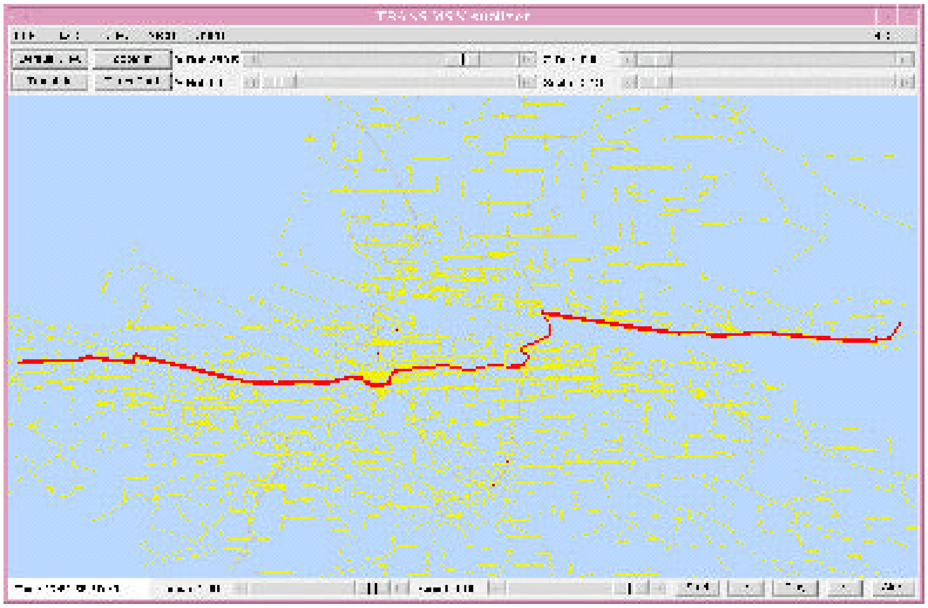

Figure 20: A single route viewed with the Output Visualizer.

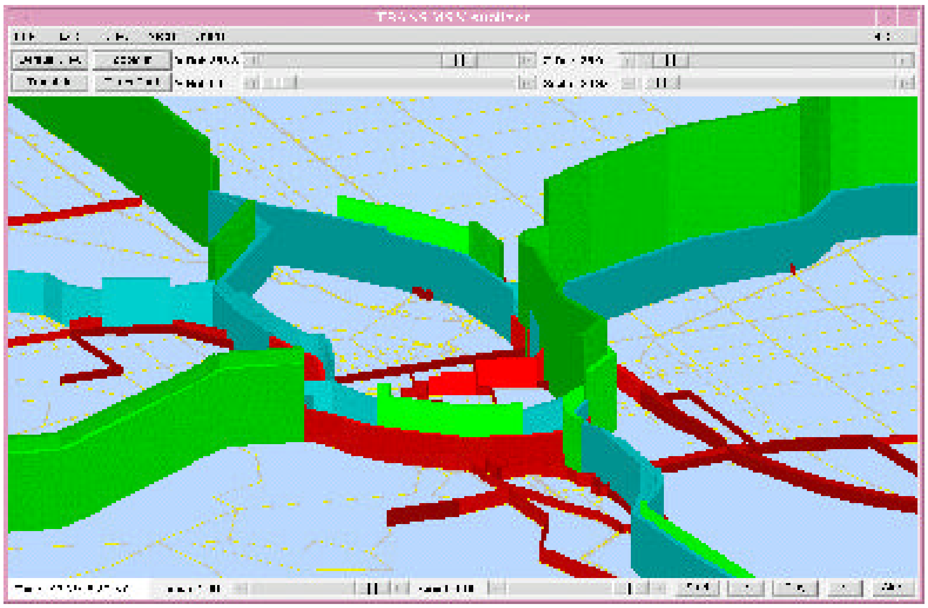

Figure 21: Multiple routes are overlaid in time to show predicted congestion points.

Traffic Microsimulator snapshot data may be displayed dynamically. Here, the individual vehicles are shown as they move on the transportation network. A picture of a single timestep for this display is shown in Figure 22. Summary data from the Traffic Microsimulator output may also be displayed by the tool. Figure 23 shows the traffic density over a short time period. Here red indicates areas where the traffic is heaviest.

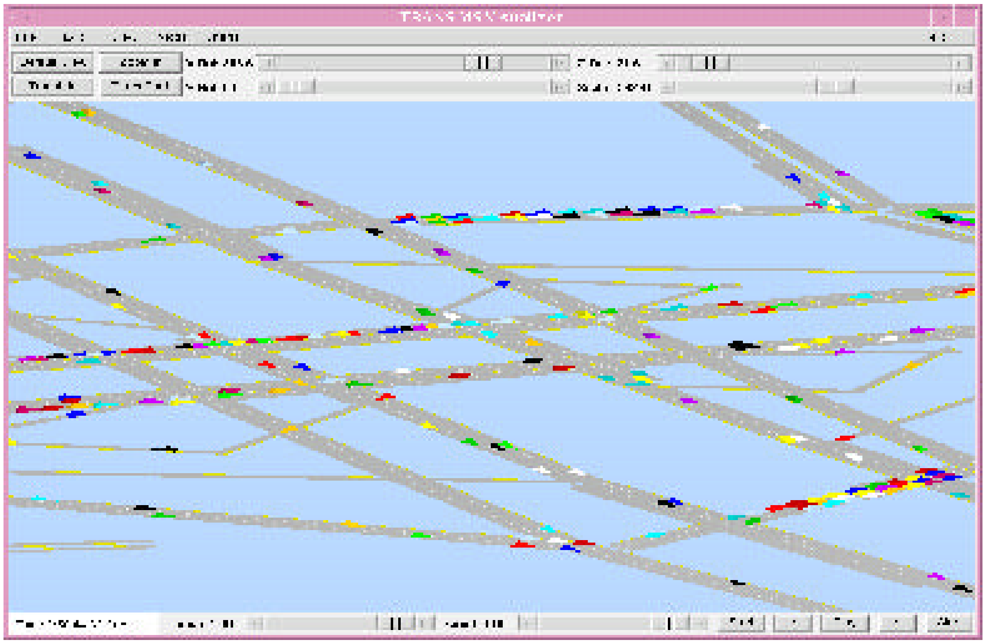

Figure 22: A one-second view of network traffic from the Traffic Microsimulator snapshot data. The Output Visualizer shows the dynamic second-to-second movement of vehicles on the roadway network.

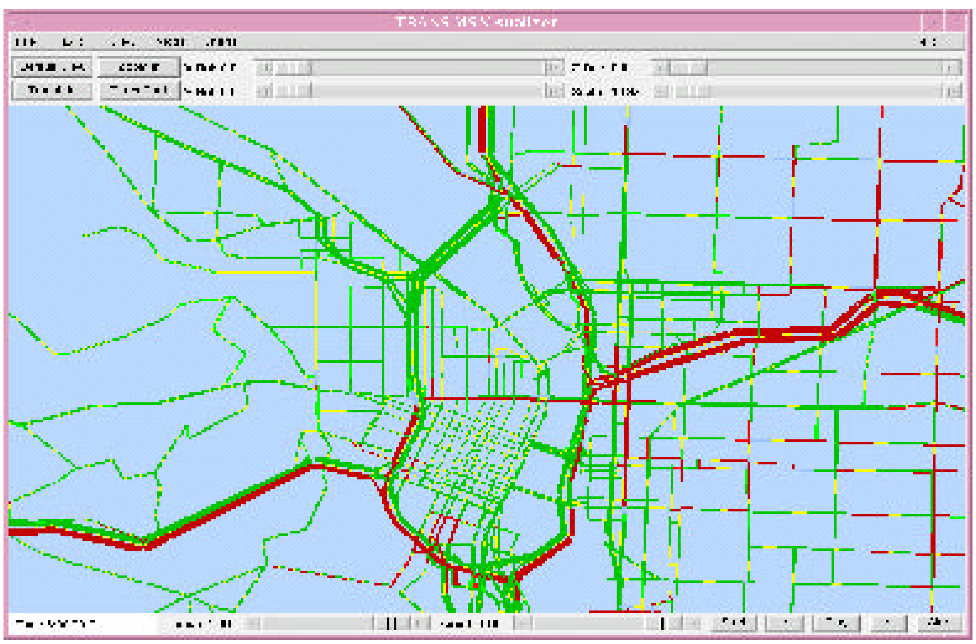

Figure 23: Traffic density from the microsimulation over a short time period is displayed. Red indicates heavily congested roadways.

5. REFERENCES

[1] Beckman, Richard J., Keith A. Baggerley, and Michael D. McKay. 1996. "Creating Synthetic Baseline Populations," *Transportation Research A*, 30 #6**,** pp. 415–429.

[2] Effa, Robert C., and Lawrence C. Larson. 1993. "Development of Real-World Driving Cycles for Estimating Facility-Specific Emissions from Light-Duty Vehicles," California Environmental Protection Agency, Air Resources Board, presented at the Air and Waste Management Association Specialty Conference on Emission Inventory, Pasadena, California, Oct. 18–20, 1993.

[3] EPA (US Environmental Protection Agency). 1993. "Federal Test Procedure Review Project: Preliminary Technical Report," EPA 420-R-93-007, Office of Air and Radiation, Washington, DC.

[4] Barth, M., T. Younglove, T. Wenzel, G. Scora, F. An, M. Ross, and J. Norbeck. 1997. "Analysis of Modal Emissions for a Diverse in-use Vehicle Fleet," *Transportation Research Record* 1587, pp. 73-84.

6. ADDITIONAL BIBLIOGRAPHY

Barrett, C., R. Jacob, and M.V. Marathe. "Formal Language Constrained Path Problems," in Proc. 6th Scandinavian Workshop on Algorithmic Theory (SWAT), Stockholm, Sweden, LNCS 1432, pp. 234-245, Springer Verlag, New York (1998).

Jacob, R., M.V. Marathe, and K. Nagel. "A Computational Study of Routing Algorithms for Realistic Transportation Networks," 2nd Workshop on Algorithmic Engineering, Saarbrucken, Germany, August 1998, Los Alamos National Laboratory report LA-UR-98-2249 (May 1998).

Jacob, R., C. Barrett, and M.V. Marathe. "Models and Algorithms for Routing in Inter-modal Time Dependent Transportation Networks," Los Alamos National Laboratory, in preparation (January 1999).

Nagel, K., and M. Schreckenberg. "A Cellular Automaton Model for Freeway Traffic," *J. Phys. I. France* **2**, pg. 2221 (1992).

Nagel, K., and M. Schreckenberg. "Traffic Jam Dynamics in Stochastic Cellular Automata," Los Alamos National Laboratory report LA-UR-95-2132 (June 1995).

Nagel, K., C. L. Barrett, and M. Rickert. "Parallel Traffic Microsimulation by Cellular Automata and Application for Large Scale Transportation Modeling," Los Alamos National Laboratory report LA-UR-96-0050 (January 1996).

Nagel, K., M. Rickert, and C. L. Barrett. *Large Scale Traffic Simulations, Lecture Notes in Computer Science*, ed. JMLM Palma and J. Dongarra, vol. 1215, pp. 380–402, Springer Verlag, New York (1997).

Wagner, P., K. Nagel, and D. Wolf. "Realistic Multi-Lane Traffic Rules for Cellular Automata," Los Alamos National Laboratory report LA-UR-96-2586 (July 1996).

Nagel, K., P. Stretz, M. Pieck, S. Leckey, R. Donnelly, and C. L. Barrett. "TRANSIMS Traffic Flow Characteristics," Los Alamos National Laboratory report LA-UR-97-3530 (July 1997).